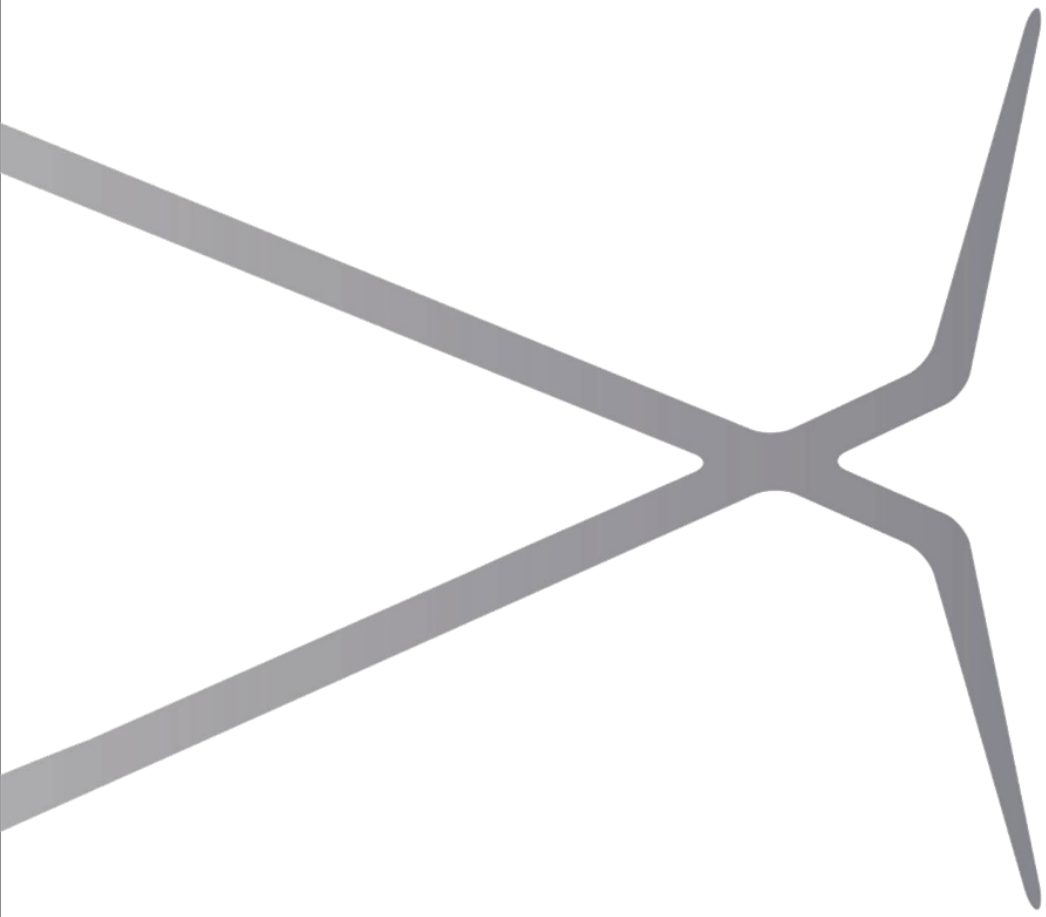

bullx

instruments for innovation

powered by Bull

Table of Contents

- HPC Overview.
- Cluster Overview.
- FLOPS.

HPC Overview

Ares, Gerardo, HPC Team

Architect of an Open World™

HPC concepts

- **HPC: High Performance Computing.**
- **Motivation:**
  - Advances in technologies have helped to address bigger complex problems.
- **In recent years were developed high performance infrastructure for scientific calculus.**
- **In the past 10 years clusters architecture has grown from 2.2% to 84.8% (source http://www.top500.com).**

Architecture	June 2000	June2010
Constelations	18.6	0.4
Clusters	**2.2**	**84.8**
MPP	51.4	14.8
SMP	27.8	0.0

HPC concepts

- **Sastify the growing requisites of the computing power.**
  - Complex problems.
  - Complex models.
  - Huge data sets.
  - Responses are time-limited.

- **Parallel processesing.**
  - Many process work together to resolve a common problem.
  - Domain decomposition or Functional parallelism are used to reduce the computing time of a resolution.

HPC concepts

- **Domain decomposition.**
  - Many simulations in science and engineering work with a simplified picture of reality in which a *computational domain*, e.g., some volume of a fluid, is represented as a grid that defines discrete positions for the physical quantities under consideration.
  - The goal of the simulation is usually the computation of observables on this grid.
  - A straightforward way to distribute the work involved across workers, i.e. processors, is to assign a part of the grid to each worker.
  - This is called *domain decomposition.*

HPC concepts

- **Functional decomposition.**
  - Sometimes, solving a complete problem can be split into more or less disjoint subtasks that may have to be executed in some specific order, each one potentially using results of the previous one as input or being completely unrelated up to some point.
  - The tasks can be worked on in parallel, using appropriate amounts of resources so that load imbalance is kept under control.

Cluster concepts

- **Parallel programming:**
  - Processing power (SMP systems).
  - Net (data communicatons).
  - Libraries development and programming API.
- **Shared memory processing:**
  - Multithread programming.
- **Distributed processing:**
  - High performance networks.
  - Independent machines interconnected by a high performance network.
  - Sincronization:
    - Available by message passing mechanism.

Cluster Overview

Ares, Gerardo, HPC Team

Architect of an Open World

Cluster overview

- **A Cluster is composed by:**
  - Compute nodes.
  - Administration nodes: login, administration.
  - I/O nodes.
  - Networks:
    - Application network.
    - Administration network.
    - I/O network.
- **Software components:**
  - Operating system.
  - Compilers.
  - Scientific and parallel libraries.
  - Management software.
  - Resource Manager.

Cluster overview

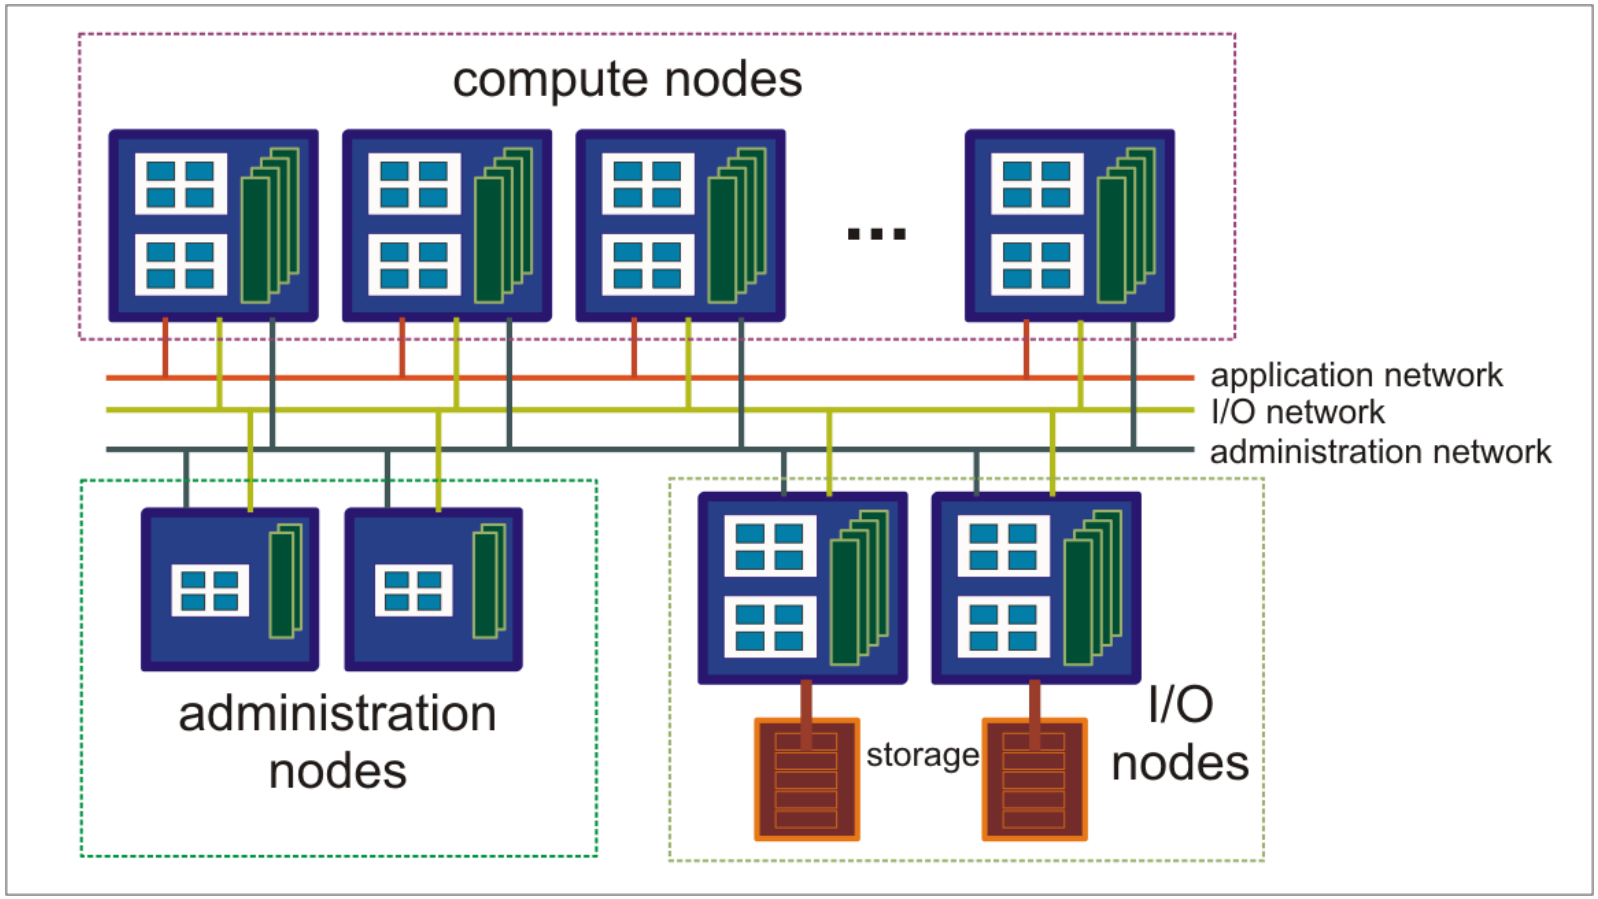

UFMG Cluster

- **UFMG cluster:**
  - Compute nodes: veredas[2-107]: 106 nodes.
    - 2 x Intel Xeon X5355 2.66GHz (4 cores).
    - 16 GB RAM memory.
  - Administration nodes: veredas0.
    - Login node.
    - Cluster administration.
  - I/O nodes: veredas[0-1].
    - NFS.
  - Networks:
    - Application network: Infiniband.
    - Administration network: 1 GbEthernet.
    - I/O network: 1GbEthernet.

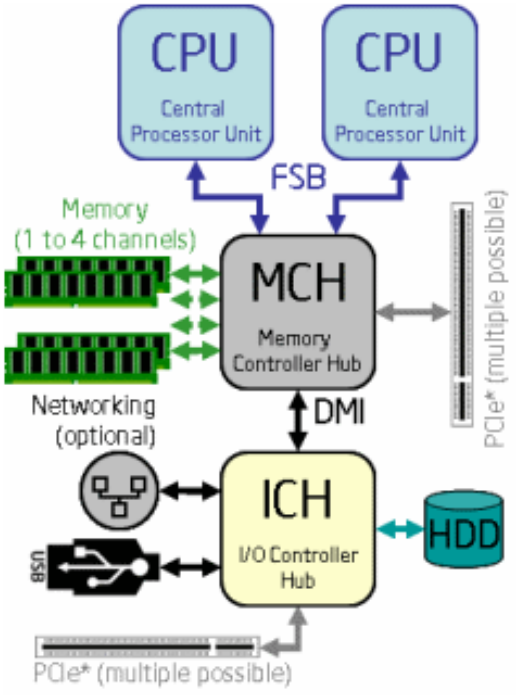

UFMG Cluster

- **Software components:**
  - Operating system: RedHat Enterprise Linux 5.3.
  - Compilers:
    - Intel C/C++ & Fortran.
    - GNU gcc & g77.
  - Scientific libraries:
    - BLACS, LAPACK, SCALAPACK, FFTW, NETCDF, HDF5, etc.
  - Parallel libraries:
    - Bull MPI 2.
    - Intel MPI.
  - Management Software:
    - Bull XBAS 5v3.1u1.
  - Resource Manager:
    - SLURM.

Π demo

Π demo

FLOPS

Ares, Gerardo, HPC Team

Flops

- One of the most accepted metrics to evaluate a cluster power is the number of Flops that it could reach.
- FLOPS is an acronym: FLoating point OPerations per Second.
- A cluster has a theorical peak of number of FLOPS in double precision.
- The Intel 5000 family series has a 128 bit SSE register.
- Each core can run two float operation in a clock tick.
- As the SSE register is 128 bit long, each core can run 4 double precision float operations.
- On an Intel Xeon X5355 Quad-core processor can be done 16 double precision float operations on each clock tick.

UFMG Flops

- To get the number of floating point operations per second on a Intel X5355 processor:

#Flops = 4 * CPU_speed * number_of_cores)

#Flops = 4 * 2.66 * 4 = 42,56

- The UFMG cluster has 106 compute node with 2 Intel X5355 processors, so the theoritical peak of the cluster is:

Theoritical peak = 106 * 2 * 42,56 = 9022,7 Gflops = 9,0 Tflops.

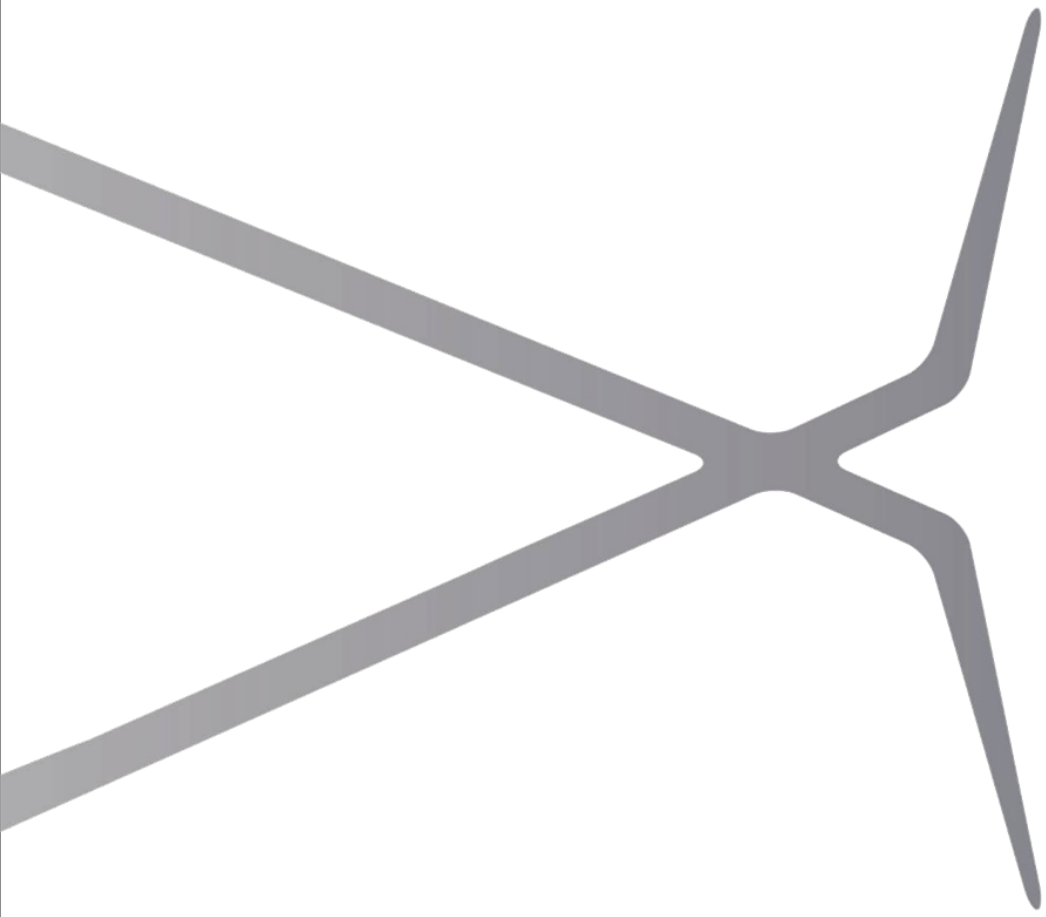

bullx

instruments for innovation

powered by Bull