



Centro di Competenza sul Calcolo Scientifico

OCCAM: a Flexible, Multi-purpose and Extendable HPC Cluster

Marco Aldinucci, Paolo Pasteris,
Sergio Rabellino

*Department of Computer Science and C3S,
University of Torino*

Stefano Bagnasco, Stefano Lusso

*Istituto Nazionale di Fisica Nucleare, sezione di Torino
and C3S, University of Torino*



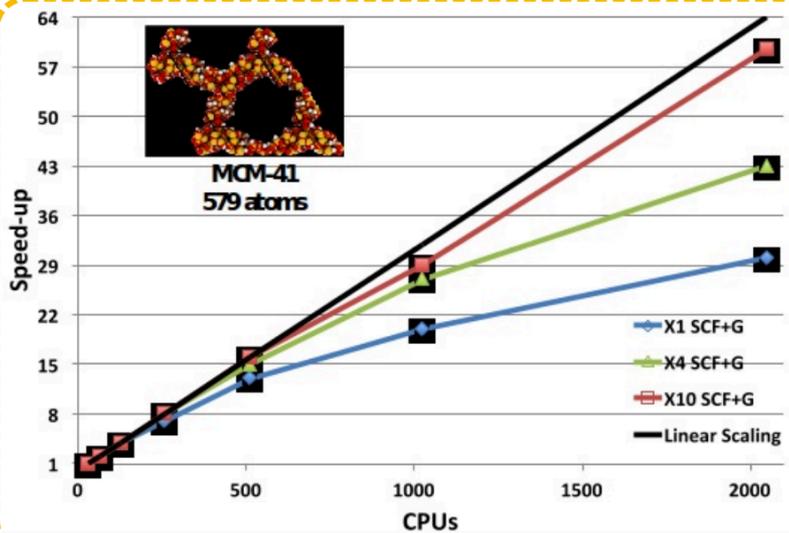
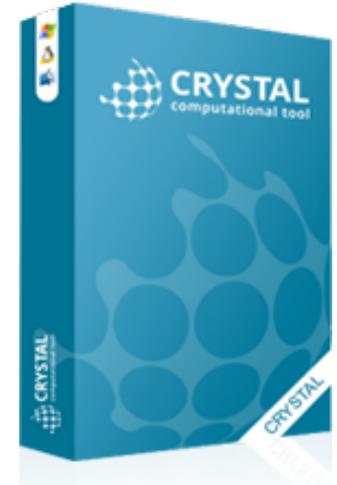


The case for an Open Computing Cluster: the use-cases

- A **very large array** of scientific use-cases from 18 university departments:
 - Computational chemistry
 - Genomics, transcriptomics & other -omics
 - Complex systems in several disciplines
 - HEP (and more) code testing & porting
 - Pharmacology & drug discovery
 - Big Data in economics & the social sciences
 - ...you name it, we have it.

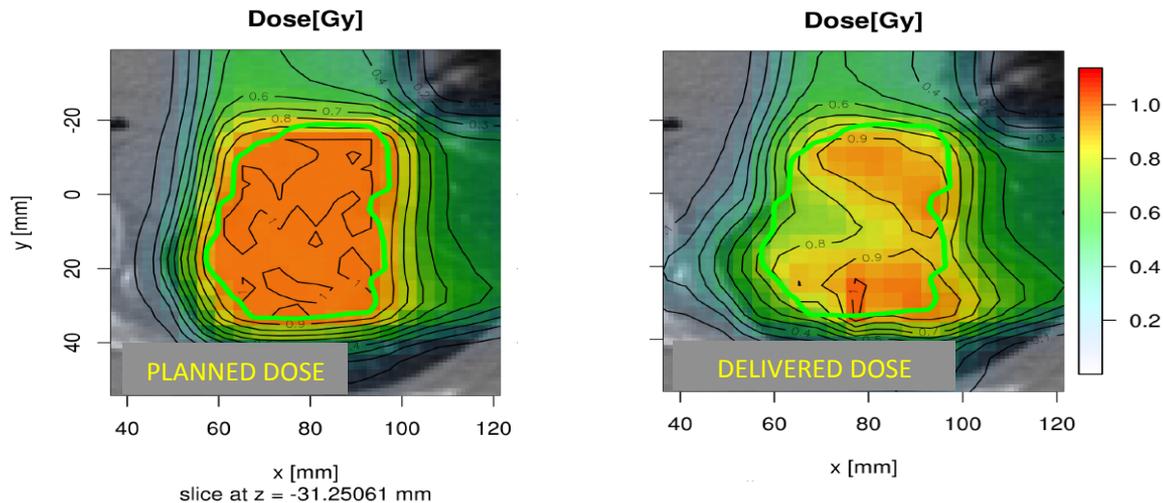
- Computational chemistry

- Software developed in Torino since the '70, used worldwide for QM computation
- Departments of Chemistry, Earth Sciences and the Centre for Nanostructured Interfaces and Surfaces



CRYSTAL can be efficiently run both in parallel (P) and massive parallel (MPP) mode. The MPP mode is particularly suitable for studying large systems, drastically improving the speed-up and reducing the requirement for memory.

- Medical physics: adrotherapy dose profiling
 - Pre-treatment calculation of reference dose (treatment planning) not identical to actual delivered dose
 - GPU-accelerated processing can allow for real-time dose profile monitoring



- Genomics, epigenomics, transcriptomics, metabolomics, proteomics, ...
 - Some use-cases (e.g. genome alignment) are highly parallel and can be run on HPC clusters
 - Others (e.g. variant calling) can benefit from very large single-image, large-memory machines



The case for an Open Computing Cluster: bridging the gap

- Very large HPC infrastructures (like CINECA in Italy)
 - In most of the cases providers of HPC services would not allow changes on the hardware or even software configuration, implementation of special features or a lower-level control on the computing infrastructure and networks
 - And often it would be useful to test code and applications on a smaller, more friendly scale before deploying on very large supercomputers...
 - ...so that researchers can have close interaction with computing experts managing the infrastructure.



The case for an Open Computing Cluster: bridging the gap

- Very small departmental clusters, with uncertain funding and understaffed maintenance
 - Uncertain and sparse funding
 - Pooling across departments never happened
 - Lack of professional support, best-effort work from researcher themselves
 - Researchers are (obviously) interested in their field, not in scientific computing itself (so they don't do R&D and are usually very conservative with the tools and technologies they use)



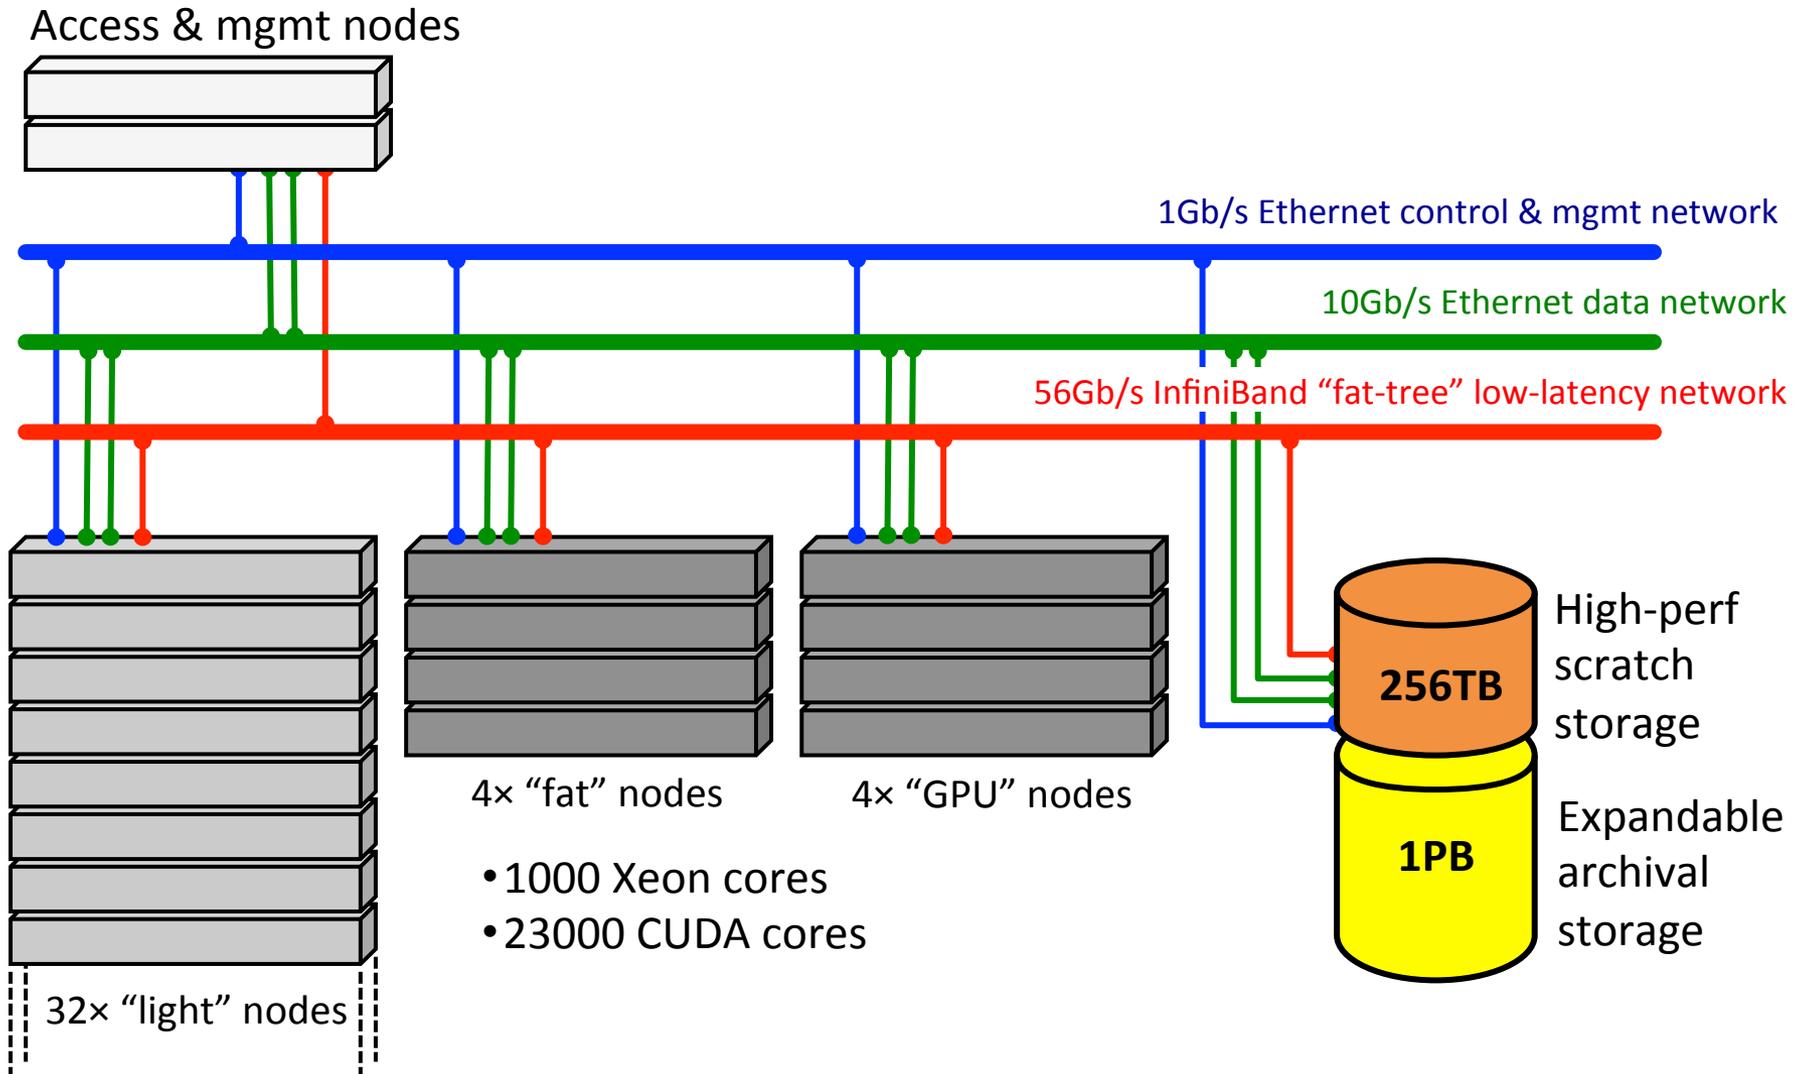
The case for an Open Computing Cluster: a platform for R&D

- Need for a platform for Scientific Computing R&D:
 - new hardware architectures (GPUs, many-core, FPGA,...)
 - scheduling & workload management
 - software technologies
 - code porting & optimisation
 - PaaS technologies
 - ...

The Open Cluster will be...

- ...highly **flexible** and **reconfigurable**:
 - conventional HPC farm, large single-image machines, accelerator-enabled platform,...
 - availability of both virtualized and bare-metal environments, dynamically partitionable
 - close interaction between the technical management and scientific users
- ... a platform for **R&D** and **innovation**:
 - Management model more like a collaborative research center than a service-oriented facility
 - Multi-disciplinary point of view will be an asset to promote new technologies, both hardware and software, in scientific computing practices
 - Will foster collaboration between Computer Science experts, technology providers, developers and scientific users (see later)

The architecture



Prospective collaboration examples

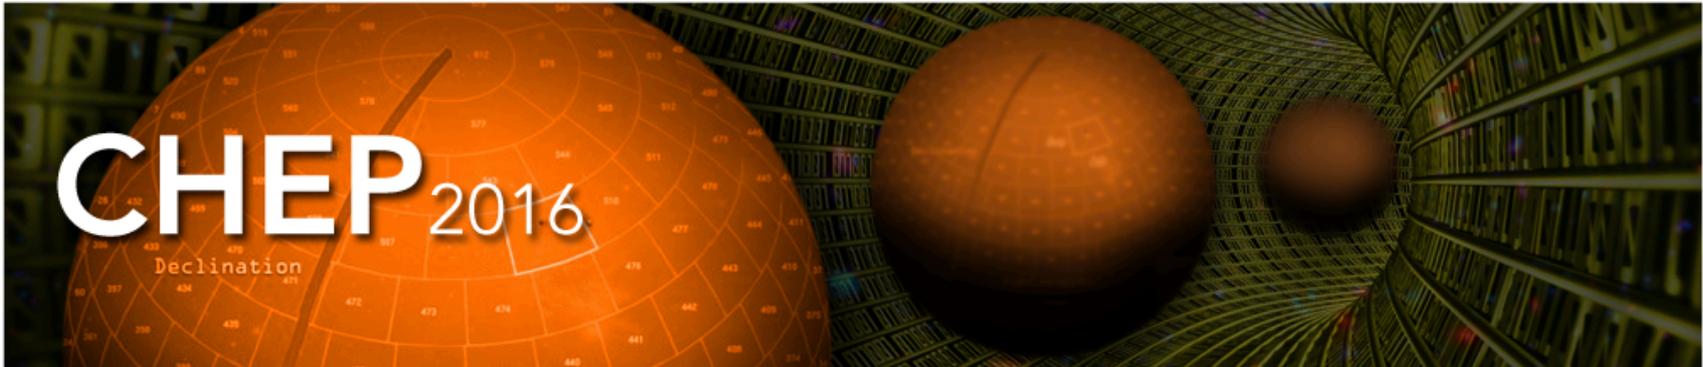
- The Computer Science Department at UNITO hosts an nVidia GPU Research Center, for example investigating the porting and optimization of HEP MC code on GPU platforms.



- The Computer Center staff at INFN-Torino is part of the INDIGO-DataCloud EU project, developing tools for PaaS services for Scientific Computing



- We designed and built an **Open Computing Cluster** that will be both a **service provider** for multi-disciplinary research and a **platform for R&D** in Scientific Computing technologies
- Also, this will serve as a laboratory to foster **collaboration** between Computer Science experts, technology providers, developers and scientific users



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

The OCCAM cluster and the Centro di Competenza sul Calcolo Scientifico of the University of Torino were funded through a contribution by the Compagnia di San Paolo



32 “Light” nodes

- CPU - 2x Intel® Xeon® Processor E5-2680 v3, 12 core 2.5Ghz
- RAM - 128GB/2133 (8 x 16 Gb)
- DISK - SSD 400GB SATA 1.8 inch.
- NET - IB 56Gb + 2x10Gb
- High density form factor (4 nodes x RU)

4 “Fat” nodes

- CPU - 4x Intel® Xeon® Processor E7-4830 v3 12 core/2.1Ghz
- RAM - 768GB/1666MHz (48 x 16Gb) DDR4
- DISK - 1 SSD 800GB + 1 HDD 2TB 7200rpm
- NET - IB 56Gb + 2x10Gb

4 “GPU” nodes

- CPU - 2x Intel® Xeon® Processor E5-2680 v3, 12 core 2.1Ghz
- RAM - 128GB/2133 (8 x 16Gb) DDR4
- DISK - 1 x SSD 800GB sas 6 Gbps 2.5”
- NET - IB 56Gb + 2x10Gb
- GPU - 2 x NVIDIA K40 su PCI-E Gen3 x16

High-performance “Scratch” storage

- DISK TYPE - HDD da 4 TB SAS 7200 rpm
- CAPACITY - 320 TB RAW e 256 TB usable
- NET - 2 x IB 56Gb FDR + 2 x 10Gb
- FILESYSTEM - Lustre Parallel Filesystem

“Archival” (non-custodial) storage

- DISK TYPE - 180 x 6 TB a 7200 rpm SAS 6Gbps
- CAPACITY - 1080 TB raw (768 TB usable)
- NET - 2 x IB 56Gb + 4 x 10GbE
- FILESYSTEM - NFS export
- Dynamic Disk Pools equivalent to RAID 6

Networking:

- InfiniBand layer - 56 Gbps “Fat Tree”
- 10GBPS Ethernet - 10 Gbps flat
- 1GBPS Ethernet for monitoring and management