Towards a data science infrastructure: Cineca's Position in the European Supercomputing Infrastructure

G. Erbacci, G. Fiameni, (*Cineca, Italy*)

Cineca, established in 1969, is a non-profit consortium of 69 Italian Universities, under the control of the Ministry of Education, University and Research. Cineca is currently the largest supercomputing centre in Italy providing cutting-edge technology and support to both academic and industry research. The mission of Cineca is to make the Italian and European research community to accelerate the scientific discovery using HPC resources in a profitable way, exploiting the newest technological advances in HPC, data management, storage systems, tools, services and expertise at large.

Cineca also represents Italy in PRACE as one of the four Tier-0 hosting centres. It is also one of the founding members of the European Technology Platform for HPC (ETP4HPC). Besides, Cineca is the Italian representative in the pan-European EUDAT Collaborative Data e-Infrastructure and one of the partner providing HPC infrastructure to the Human Brain Project, the EU flagship project facing the big challenge of understanding the human brain.

The current CINECA HPC infrastructure is composed of:

- **Tier-0 FERMI**: IBM BlueGene/Q (10 Frames, 163840 cores with 1GB RAM per core and a peak performance of 2.1 PFlop/s), Ranked No 23 in the Top500 (November 2014);
- **Tier-1 Galileo**: Hybrid Linux cluster (8192 cores Intel Haswell with 26 TBytes RAM and 760 Intel Xeon Phi 7120P accelerators, with a peak performance of 1 PFlops);
- **PICO**: the High Performance Data Analytics Linux Cluster (80 nodes Intel NextScale) integrated with a multi tier storage system configured with 40 TBytes of SSD memory, 5 PBytes of High IOPS storage and 10 PBytes of long term archive.

Our belief is that the new data oriented infrastructure would become a valuable asset for Italian and European scientists while supporting an open science policy. Optimizing the access to data is also a key factor to enable multi-disciplinary research thus reducing the duplication of effort and cost while accelerating scientific progress.

With the introduction of PICO, the improvement of the storage capacity and capabilities, and the reorganization of the network connectivity, the traditional Cineca compute-centric HPC infrastructure has made a step forward the realization of a Data Centric infrastructure, as represented in Figure 1, where data cover a fundamental role in any scientific application. In our view, the transition from a silos centric organization where different computing platforms rarely interact together towards a more integrated, data oriented environment, where various types of workloads are supported and data can be accessed from all computing engines efficiently, represent an important breakthrough with respect to the past.

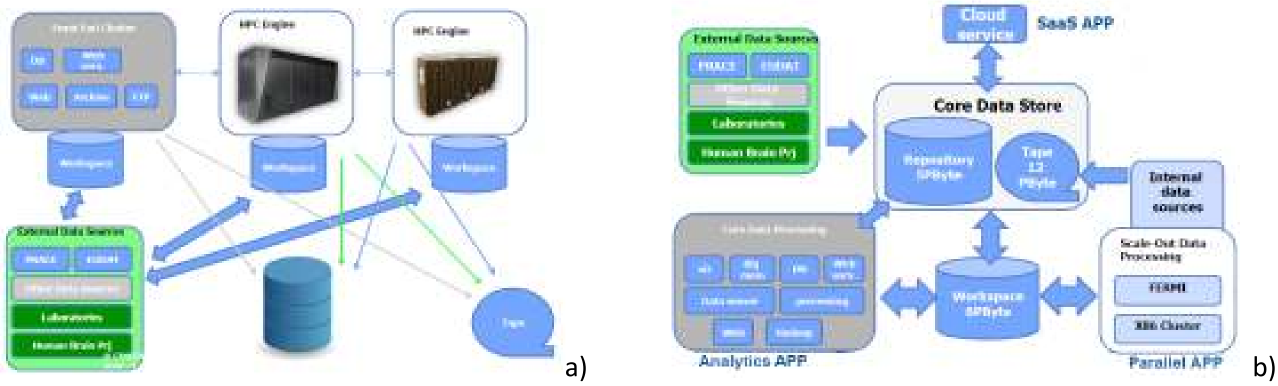

Figure 1: a) HPC island Infrastructure and b) Data centric infrastructure

The new infrastructure aims at responding to the ever-growing demand of services and capabilities (storage, management, compute, and visualization) which are essential to address the challenge of "big data" problems, nowadays a key factor of many scientific disciplines. Actually we assist to a more and more consolidation of data-intensive scientific discovery (Data Science) which impose a data centric view.

Through this approach, users can easily integrate the main HPC infrastructure (FERMI) with the data intensive computing facility (PICO). The peculiar hardware characteristics (large memory per node, massive storage equipment and sharing, fast data access and transfer, etc.), in combination with the availability of analytics frameworks, such as Apache Hadoop or Apache Spark, facilitate the execution of high-throughput applications.

The processing power combined with the access to advanced software packages enables researchers to harness the potential of supercomputing for the analysis of large volume of data and can be applied to many fields including predictive analytics, social media analytics and text analytics to disease detection, prevention and treatment, financial modeling and smart energy metering.

A new range of new services, usually uncommon within the HPC domain, can be provided, specifically in the analytics domain, including the support on data mining and statistical software and the expertise on machine learning, predictive modeling, text mining, data pre-processing and data visualization for targeted applications development.

In addition, Cineca is also involved in specific projects in the fields of epigenomic (www.epigen.it), Next Generation Sequencing (NGS), connectome-inspired computing (http://www.iconfoundation.net), as well as Human Brain Project (http://www.humanbrainproject.eu). Further projects, both in the scientific, industrial and public field will be treated as use cases for the concrete adoption of the new data-centric approach and will help defining the best way to structure future data services (e.g. as SaaS or IaaS).

With respect to the evolution of the infrastructure, our vision is to follow a two steps roadmap (2015-16 and 2017-18), which should lead us towards an integrated Tier-0 system (50 PFlops and a congruent amount of High IOPS storage, in the order of 50 PByte) coupling large scale (HPC), high throughput (HTC) and high performance data analytics (HPDA) capabilities together, as reported in Figure 2. Different layers will manage the facility, as reported in Figure 3. This vision represents a first step toward the evolution to an Exascale Infrastructure.

Logical Name	FERMI Tier-0 (June 2014)	Galileo Tier-1 (Dec 2014)	PICO Big Data (Oct 2014)
Peak Performance	~ 2.1 PFlops	~ 1 PFlops	5 PBytes High IOPS storage

Logical Name	Tier-0 (2015-2016)	BIG DATA (2016)
Peak Performance	~ 10 PFlops	~ 1 PFlops, 10 PByte High IOPS Storage

Logical Name	Tier-0 & BIG DATA (2017-2018)
Peak Performance	~ 50 PFlops, 50 PByte High IOPS Storage

Figure 2: Cineca Infrastructure Roadmap

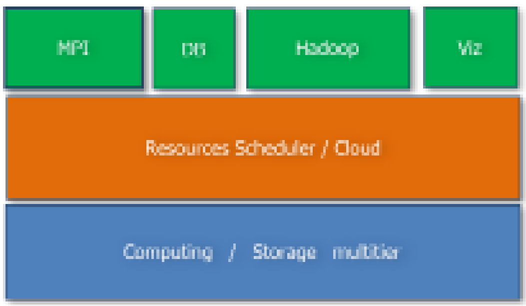

Figure 3: Cineca Infrastructure Layers

Extreme data applications integrated with HPC facility are gaining traction in science and industry and big data analytics as a whole is growing in importance by providing unprecedented and intuitive insights in science but also into an array of business functions such as sales, supply chain, and network operations.

Future large-scale simulations running on next generation, pre-exascale supercomputers will be data-intensive and require interactive analysis, visualization and control.

The Cineca pre-exascale infrastructure will investigate new paradigms, like interactive supercomputing, where large memory capacity is required, provided by a new memory hierarchy layer based on very dense and power-efficient memory technologies. Interactive supercomputing need different operation modes from the current ones. Large-scale simulation jobs, data analysis jobs, visualization pipelines, etc. will cohabit on the same system. Users will require to run multiple jobs concurrently within a single session with resources that will change dynamically.

The combination of extreme-scale data analysis and HPC applications, although of remarkable interest presents several challenges due to the change of paradigms, but also due to the volume, variety, and velocity at which data is collected, for instance, from environmental sensors. To be profitable, provided services and capabilities must be agile, quickly adapting to new needs, and to changing scientific and business requirements.

With this respect a lot of challenges exist and important key points, some listed below, must be solved to properly address this infrastructure in the direction of the Exascale:

- Novel programming abstractions are needed for achieving fine-grained control over the trade-off between the data analytics latency and the accuracy of results; data-intensive algorithms do not have the well developed high performance libraries familiar to HPC;
- Develop novel programming methodologies and runtime systems for effectively exploiting extreme computing infrastructures integrated with the analysis of big data;
- Heterogeneous, energy aware and easy-to-use programming environments must be implemented;
- Fundamental advances must be done in how to express ad-hoc analytic queries over massive data;
- Tools for performance monitoring and analysis of HPC big data systems must be enhanced;
- Agile approaches for the development of highly scalable analytics applications on high-end computing systems still must be addressed;
- New extreme data analytics applications, which are both computationally and data intensive, are still lacking and need more attention from the scientific community;
- Tools and methods for interactive analytics over big data need to be explored;
- Integrated approaches for extreme heterogeneity management supercomputers toward Exascale are needed;
- Resource requirements are expected to change during interactive job execution, and thus mechanisms to free compute resources or to use them malleably are needed;
- Implement tools for the management of dense memory, supporting malleability and job pre-emption;
- Implement check-pointing in the context of dynamic resource management.