

# Low latency network and distributed storage for next generation HPC systems: the ExaNeSt project.

*Tuesday, 11 October 2016 15:00 (15 minutes)*

With processor architecture evolution, the HPC market has undergone a paradigm shift. The adoption of low-cost, Linux-based clusters extended HPC's reach from its roots in modeling and simulation of complex physical systems to a broad range of industries, from biotechnology, cloud computing, computer analytics and big data challenges to manufacturing sectors. In this perspective, the near future HPC systems will be composed of millions of low-power-consumption computing cores, tightly interconnected by a low latency high performance network, equipped with a new distributed storage architecture, densely packaged but cooled by an appropriate technology.

In the road towards Exascale-class system, several additional challenges wait for a solution; the storage and interconnect subsystems, as well as a dense packaging technology, are three of them.

The ExaNeSt project, started on December 2015 and funded in EU H2020 research framework (call H2020-FETHPC-2014, n. 671553), is a European initiative aiming to develop the system-level interconnect, the NVM (Non-Volatile Memory) storage and the cooling infrastructure for an ARM-based Exascale-class supercomputers. The ExaNeSt Consortium combines industrial and academic research expertise, especially in the areas of system cooling and packaging, storage, interconnects, and the HPC applications that drive all of the above.

ExaNeSt will develop an in-node storage architecture, leveraging on low cost, low-power consumption NVM devices. The storage distributed sub-system will be accessed by a unified low latency interconnect enabling scalability of storage size and I/O bandwidth with the compute capacity.

The unified, low latency, RDMA enhanced network will be designed and validated using a network test-bed based on FPGA and passive copper and/or active optical channels allowing the exploration of different interconnection topologies (from low radix n-dimensional torus mesh to higher radix DragonFly topology), routing functions minimizing data traffic congestion and network support to system resiliency.

ExaNeSt also addresses packaging and advanced liquid cooling, which are of strategic importance for the design of realistic systems, and aims at an optimal integration, dense, scalable, and power efficient. In an early stage of the project an ExaNeSt system prototype, characterized by 1000+ ARM cores, will be available acting as platform demonstrator and hardware emulator.

A set of relevant ambitious applications, including HPC codes for astrophysics, nuclear physics, neural network simulation and big data, will support the co-design of the ExaNeSt system providing specifications during design phase and application benchmarks for the prototype platform.

In this talk a general overview of project motivations and objectives will be discussed and the preliminary developments status will be reported.

## Secondary Keyword (Optional)

Network systems and solutions

## Primary Keyword (Mandatory)

High performance computing

## Tertiary Keyword (Optional)

**Primary author:** VICINI, Piero (Universita e INFN, Roma I (IT))

**Presenter:** VICINI, Piero (Universita e INFN, Roma I (IT))

**Session Classification:** Track 6: Infrastructures

**Track Classification:** Track 6: Infrastructures