Contribution ID: **49** Type: **Oral Presentation**

UniNuvola: the computing portal of the Perugia University

Tuesday, 26 March 2024 11:00 (30 minutes)

An innovative, distributed, and elastic computing ecosystem, called UniNuvola, is being deployed at the University of Perugia, involving the Department of Chemistry, Biology and Biotechnologies, the Department of Physics and Geology, the Department of Mathematics and Informatics and the Department of Engineering. The aim of the project is the creation of a federated and scalable computing infrastructure, providing scientific services to end users belonging to both academic structures and, in perspective, SMEs. It represents a proof of concept of a distributed computing infrastructure empowered with software-defined networking capabilities, with every node laying transparently its virtual resources, services, and microservices on a virtual backbone. Our main objective is to design a virtual distributed networking infrastructure to federate an innovative architecture, with dynamical resource allocation, intelligent management of large volumes of data and compliant with present European federated computation paradigm and data protection policies. The federation also manages a heterogeneous collection of virtual application environments by organizing them into well tested and auto-consistent packages, ready to be used by organizations guaranteeing, at the same time, the requested performance and result accuracy features. To this end, a first prototype of four Dell Power Edge R940
nVME servers, each equipped with 2 Intel Xeon Gold CPU, 512 GB of RAM, and 16 TB of disk space, has been configured with the most recent software solutions for pursuing the above mentioned objectives. More in detail, a Kubernetes cluster has been installed, using the Rook operator for deploying a CEPH scalable distributed storage. We have also investigated the adoption of both Metallb and OVN load balancers. User authentication has been managed with the Vault server interfaced to the University LDAP, while Jupyter Hub has been containerized into Kubernetes to serve Notebooks for multiple users. On the other hand, for those workloads already running in virtual environments that are difficult to containerize into Kubernetes pods, KubeVirt technology provides us with the possibility of enabling KVM-based virtual machine workloads to be managed as pods.
Upon this (virtual) infrastructure a ready-to-use collection of scientific packages, for both research and education purposes, is being developed. To this end, the capabilities of UniNuvola will be firstly benchmarked with various use cases, built upon computational chemistry and machine learning applications. Computational chemistry applications are widely recognized for their high demands on CPUs and storage, making them the ideal candidates for testing the scalability of the architecture. The significance of machine learning lies not only in its wide range of applications but also in its high memory requirements. Applications ranging from image recognition to intrusion detection algorithms, tested on well-known datasets or from live interaction with external sources, will allow benchmarking the capability of the platform with respect to the commercial, widespread alternatives. In addition, both cases can be utilized to test future improvements to the infrastructure, such as the inclusion of GPUs and quantum computing. In fact, during the second phase of the project, both high-end nodes with GPUs (Nvidia A100) and a solid-state SpinQ Triangulum NMR quantum computer will be integrated into the UniNuvola cluster to realize an academic across-the-board data center able to serve a variety of instances.

Primary authors: Prof. FANO', Livio (Department of Physics and Geology, University of Perugia); Prof. FEMMINELLA, Mauro (Department of Engineering, University of Perugia); Prof. POGGIONI, Valentina (Department of Mathematics and Informatics, University of Perugia); Prof. BATTISTONI, Michele (Department of Engineering, University of Perugia); Dr PACIFICI, Leonardo (Department of Chemistry, Biology and Biotechnologies, University of Perugia); Prof. ANGELONI, Laura (Department of Mathematics and Informatics, University of Perugia); Prof. COSTANTE, Gabriele (Department of Engineering, University of Perugia); Dr MARIOTTI, Mirko (Department of Physics and Geology, University of Perugia); Dr PEZZELLA, Marco (Department of Physics and Geology, University of Perugia); Prof. FAGINAS LAGO, Maria Noelia (Department of Chemistry, Biology and Biotechnologies, University of Perugia); Prof. LOMBARDI, Andrea (Department of Chemistry, Biology and Biotechnologies, University of Perugia); Dr VITILLARO, Giuseppe (CNR-SCITEC)

Presenter: Dr PEZZELLA, Marco (Department of Physics and Geology, University of Perugia)

Session Classification: Converging High Performance infrastructures: Supercomputers, clouds, accelerators

Track Classification: Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators