

Cloud benchmarking and validation test suite

Monday, 2 November 2020 15:00 (15 minutes)

The Helix Nebula Science Cloud (HNSciCloud) project developed a hybrid cloud linking commercial cloud service providers and research organisations in-house resources via the GÉANT network. Following HNSciCloud, the Open Clouds for Research Environments project (OCRE) leverages its experience in the exploitation of commercial cloud services, currently being considered by the European research community as part of a hybrid cloud model to support the needs of their scientific programmes. Parallel to OCRE, the Archiving and Preservation for Research Environments (ARCHIVER) project -led by CERN- combines multiple Information and Communications technologies, including extreme data-scaling, network connectivity, service interoperability and business models, in a hybrid cloud environment to deliver end-to-end archival and preservation services that cover the full research lifecycle.

During the testing phases of HNSciCloud it became evident how challenging such tasks can be as cloud services providers offer a variety of products that are often unknown to the research community. In order to test and compare services performance and adequacy for multiple scientific domains (High Energy Physics, Life Sciences, Photon-Neutron Sciences, Astronomy), an automated testing suite to run a set of tests and benchmarks across all cloud stacks is needed. Such a framework would strongly support testing activities on both ARCHIVER and OCRE projects.

CERN IT developed a test-suite that leverages on the testing activities of HNSciCloud where the European Research organisations involved -CERN being one of them- representing multiple use cases from several scientific domains, put together more than thirty tests and benchmarks. This tool was designed to be as modular and autonomous as possible utilising open source technologies well established in industry that allow easy deployment and transparent assessment. It relies first on Terraform for resource provisioning, followed by Ansible for compute instances configuration and bootstrapping of a Kubernetes cluster, which provides the abstraction layer so eventually tests run on Docker containers. Finally, results are pushed to an S3 bucket that the CERN OpenStack cloud hosts. The test catalog offers functional and performance benchmarks in several technical domains such as compute, storage, HPC, GPUs, network connectivity performance, and advanced containerised cloud application deployments but also covers upper levels in the stack as for example evaluating the degree of "FAIRness" of data repository services or federated AAI protocols. The concept is expected to be expanded and considered a best practice for the on-boarding of commercial services for research environments in the context of the European Open Science Cloud, to ensure that cloud offerings conform to the requirements and satisfy the needs of the research community.

On this demonstration the test-suite will be run showing the whole workflow, from configuration of the tool to results gathering, using as test deployment examples a variety of networking tests and CPU benchmarks, as these two areas are relevant for most domains.

Primary authors: PELUAGA, Ignacio (CERN); URBAN, Jakub (CERN); DEVOUASSOUX, Marion; MANU, Anna (CERN); JONES, Bob (CERN); FERNANDES, Joao (CERN); JAMES, Jimmy (CERN)

Presenter: PELUAGA, Ignacio (CERN)

Session Classification: Demos 1