



Contribution ID: 258

Type: **Oral**

Running HTC and HPC applications opportunistically across private, academic and public clouds

Thursday, 7 November 2019 15:00 (15 minutes)

Access to both High Throughput Computing (HTC) and High Performance Computing (HPC) facilities is vitally important to the fusion community, not only for plasma modelling but also for advanced engineering and design, materials research, rendering, uncertainty quantification and advanced data analytics for engineering operations. The computing requirements are expected to increase as the community prepares for ITER, the next generation facility. Moving to a decentralised computing model is vital for future ITER analysis where no single site will have sufficient resource to run all necessary workflows.

The Fusion Science Demonstrator in the European Open Science Cloud for Research Pilot Project (EOSCpilot) aimed to demonstrate that the fusion community can make use of distributed cloud resources. PROMINENCE is a platform initially developed within this Science Demonstrator and enables users to transparently exploit idle cloud resources for running scientific workloads. In addition to standard HTC jobs, HPC jobs such as multi-node MPI are supported. All jobs are run in containers to ensure they will reliably run anywhere and are reproducible. Cloud infrastructure is invisible to users, as all provisioning, including extensive failure handling, is completely automated. On-premises cloud resources can be utilised and at times of peak demand burst onto external clouds. In addition to the traditional “cloud-bursting” onto a single cloud, PROMINENCE allows for bursting across many clouds in a hierarchical manner, for example bursting from a local private cloud to national research clouds, then across many clouds in the EGI FedCloud federation, and finally to public clouds. Job requirements are also taken into account, so jobs with special requirements, e.g. high memory or access to GPUs, are sent only to appropriate clouds. Several different storage options are available to either allow data to be staged-in and out of jobs using Swift/S3 or to provide POSIX-like access to data, irrespective of where the job is running.

In this presentation we will describe PROMINENCE, its architecture and the challenges of using many clouds opportunistically. We will also report on our experiences with several fusion use cases.

Consider for promotion

Yes

Primary authors: Mr DE WITT, Shaun (UKAEA); Dr CABALLER, Miguel (UPVLC); LA ROCCA, Giuseppe (EGI); Dr COSTER, David (Max-Planck-Institut für Plasmaphysik); Dr LAHIFF, Andrew

Presenter: Dr LAHIFF, Andrew

Session Classification: Track 7 –Facilities, Clouds and Containers

Track Classification: Track 7 –Facilities, Clouds and Containers