

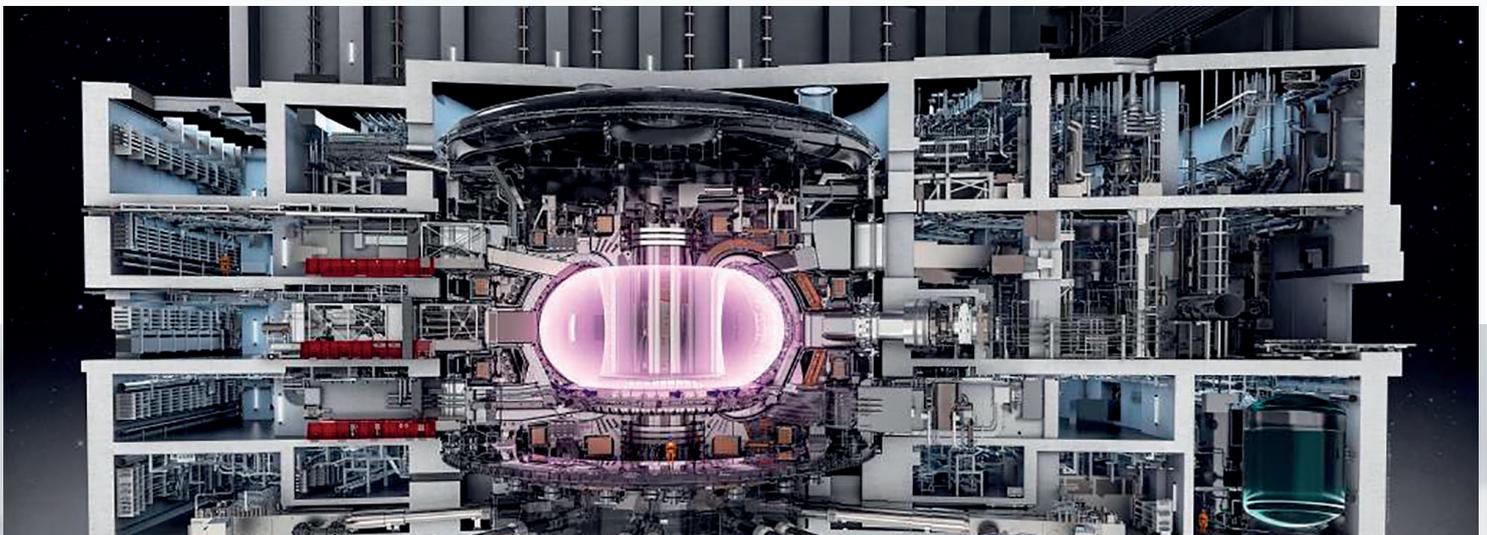
Brief Overview

Access to HPC facilities are vitally important to the fusion community, not only for plasma modelling but also for advanced engineering and design, materials research, uncertainty quantification and advanced data analytics for engineering operations (e.g. condition monitoring). The requirements for HPC class machines is expected to only increase as the community prepares for the next generation facility, ITER. However, access to HPC class infrastructure is quite restricted and obtaining time on this class of machine for algorithm development, testing and optimisation is already problematic. A few fusion centres have local access to smaller HPC class clusters but larger scale workflows and smaller fusion research centres are required to competitively bid for time on hardware such as PRACE Tier-0 and Tier-1 facilities, often also requiring visiting the centre. Within this demonstrator we propose to make HPC class machines available as a cloud like service to the fusion community.

Objectives

PROMINENCE provides smaller fusion sites access to HPC class nodes through a cloud interface. These sites are unique, e.g. in the development of novel algorithms which will lead to increased efficiency and impact. The demonstrator supports the fusion community to make use of distributed cloud resource, and ensure that sites running small HPC clusters can utilise them efficiently.

Beyond fusion, the availability of HPCaaS will be of interest to other disciplines and will be complementary to larger machines available in PRACE. Making hardware available through EOSC infrastructure with the capability of bursting between different providers is a desirable feature for communities wanting to make limited use of HPC for modelling, analysis or development and scalability testing of novel codes.





EOSC pilot

The European Open Science
Cloud for Research Pilot Project

Main achievements

- » Successfully used Indigo Infrastructure Manager (IM) to create small compute clusters on EGI FedCloud Resources
- » Demonstrated the running small scale MPI applications on those clusters to assess the impact of high latency interconnects
- » Demonstrated the ability to bust between fedcloud sites and commercial sites
- » Performance comparison of EGI FedCloud sites with commercial sites
- » Discovered currently IM does not integrate with SLA Manager
- » Demonstrated successful transfer of data from EGI FedCloud sites to S3 long term storage provided by STFC, and users ability to then access results
- » Trained a subset of users on the use of the interface and assisted them in deriving science results.

Recommendations for the implementation

Currently the need to have a virtual organization to make of EGI FedCloud resources, together with the limited amount of resources, mean we cannot promote this as widely as we would like. Better integration of EGI and Indigo services should be considered a high priority. On our side, we are finalizing a simplification of the user interface to hide the complexity which will make cloud submission look more like batch submission.

Partners of the SD

Science Partners: Chalmers University of Technology, IPP Garsching

INFRASTRUCTURE Partners: EGI (list sites like CESNET), Indigo (mention partners like INFN)



// The science demonstrator has shown that for certain classes of HPC codes, use of cloud infrastructure with its inherent scalability, can be used to supplement local resources and ensure a faster time for science delivery. //

Contacts

✉ shaun.de-witt@ukaea.uk