Proposal for a BDEC2 Platform Demonstrator from CERN and SKA

We would like to propose a proof-of-concept container platform and batch integration for workloads submission to access HPC testbed resources for Data Intensive Science applications such as HEP (LHC experiments) and radio-astronomy (SKA).
Representatives workloads, which in HEP concern the so-called “offline processing” i.e. the asynchronous data processing workflows in WLCG [1] - reconstruction and simulation - which happen after the primary raw data has been collected, would include (in order of complexity):

1. A standalone containerized benchmarking suite, that we propose to measure the performance of compute resources for our workflows; this benchmarking suite contains representative applications by each experiment for both High Throughput Computing (HTC) and High Performance Computing (HPC); it is light-weight, extensible and can be very easily modified to add other applications, including ML specific ones, and is reproducible, portable, with no need for network connectivity.

2. By leveraging the ‘OCI Image spec’, this demonstrator will show how optimised container images can be distributed to get the most performance out of the different target systems.

These workflows are available and ready to be used for the demonstrator. At a later stage we would like to extend the work to include more complex workflows, such as:

3. Fast simulation for LHC Monte Carlo detector simulation; no input data is needed. Generative models approaches are used, in particular with 3D conditional Generative Adversarial Networks (GAN); here significant speed-up for the training process can be achieved in an MPI based distributed parallel approach. On the SKA side approaches based on graph-based analytics using DASK would also be of particular interest.

4. Once workflows 1 and 2 are demonstrated, a complementary and fully IO intensive workload could be also exercised, e.g. via a full reconstruction workflow of an LHC experiment, which requires accessing large amounts of raw data, significant local IO and significant data export to a secondary analysis data format, typical of a data intensive science, like HEP and the distributed predict I/O requirements of the SKA[2].

In terms of “platform vision” and shared cyberinfrastructure (CI) we would like to demonstrate how to integrate HPC facilities with our existing world-wide distributed computing facilities, which in a way are representative of distributed computing models such as found in both the

SKA Regional Centres and SKA-SDP [3] approaches which may involve the exploitation of more agile, cloud-like platforms to accommodate a range of different execution environments. We have considerable experience in integrating our applications within external resources such as public and private clouds. The LHC experiments have also integrated “ad hoc” into several HPC centres, in particular in the USA and Europe.

We believe that this demonstrator would add significantly in terms of a harmonized and common approach to integrating HPC facilities. We can perform the integration of this demonstrator leveraging existing human resources and using our existing distributed computing software. We would be looking for partners at HPC sites, in particular to access their, as well as potentially share, test-bed resources, sufficient at this stage for a proof-of-concept.

[1]
http://cds.cern.ch/record/1695401?ln=en

[2]
http://ska-sdp.org/sites/default/files/attachments/distributed_predict_io_prototype_part_1_-_signed.pdf
[3]
http://ska-sdp.org/sites/default/files/attachments/ska-tel-sdp-0000151_01_sdp_p3alaskareport_part_1_-_signed_2.pdf