

EGI federated platforms supporting accelerated computing

Thursday, 9 March 2017 16:30 (30 minutes)

While accelerated computing instances providing access to NVIDIA GPUs are already available since a couple of years in commercial public clouds like Amazon EC2, the EGI Federated Cloud has put in production its first OpenStack-based site providing GPU-equipped instances at the end of 2015. However, many EGI sites which are providing GPUs or MIC co-processors to enable high performance processing are not directly supported yet in a federated manner by the EGI HTC and Cloud platforms. In fact, to use the accelerator cards capabilities available at resource centre level, users must directly interact with the local provider to get information about the type of resources and software libraries available, and which submission queues must be used to submit accelerated computing workloads.

EU-funded project EGI-Engage since March 2015 has worked to implement the support to accelerated computing on both its HTC and Cloud platforms addressing two levels: the information system, based on the OGF GLUE standard, and the middleware. By developing a common extension of the information system structure, it was possible to expose the correct information about the accelerated computing technologies available, both software and hardware, at site level. Accelerator capabilities can now be published uniformly, so that users can extract all the information directly from the information system without interacting with the sites, and easily use resources provided by multiple sites. On the other hand, HTC and Cloud middleware support for accelerator cards has been extended, where needed, in order to provide a transparent and uniform way to allocate these resources together with CPU cores efficiently to the users.

In this paper we describe the solution developed for enabling accelerated computing support in the CREAM Computing Element for the most popular batch systems and, for what concerns the information system, the new objects and attributed proposed for implementation in the version 2.1 of the GLUE schema.

For what concerns the Cloud platform, we describe the solutions implemented to enable GPU virtualization on KVM hypervisor via PCI passthrough technology on both OpenStack and OpenNebula based IaaS cloud sites, which are now part of the EGI Federated Cloud offer, and the latest developments about GPU direct access through LXD container technology as a replacement of KVM hypervisor.

Moreover, we showcase a number of applications and best practices implemented by the structural biology and biodiversity scientific user communities that already started to use the first accelerated computing resources made available through the EGI HTC and Cloud platforms.

Summary

In this paper we describe the solution developed for enabling accelerated computing support in the CREAM Computing Element for the most popular batch systems and, for what concerns the information system, the new objects and attributed proposed for implementation in the version 2.1 of the GLUE schema. We describe also the solutions implemented to enable GPU virtualization on KVM hypervisor via PCI passthrough technology on both OpenStack and OpenNebula based IaaS cloud sites, which are now part of the EGI Federated Cloud offer, and the latest developments about GPU direct access through LXD container technology as a replacement of KVM hypervisor.

Moreover, we showcase a number of applications and best practices implemented by the structural biology and biodiversity scientific user communities that already started to use the first accelerated computing resources made available through the EGI HTC and Cloud platforms.

Primary authors: Dr REBATTO, David (Istituto Nazionale di Fisica Nucleare - Sez. di Milano, Italy); Dr ASTALOS, Jan (Institute of Informatics Slovak Academy of Sciences, Slovakia); Dr ZANGRANDO, Lisa (Istituto Nazionale di Fisica Nucleare - Sez. di Padova, Italy); Dr VERLATO, Marco (Istituto Nazionale di Fisica Nucleare - Sez. di Padova, Italy); Dr DOBRUCKY, Miroslav (Institute of Informatics Slovak Academy of Sciences, Slovakia); Dr ANDREETTO, Paolo (Istituto Nazionale di Fisica Nucleare - Sez. di Padova, Italy); Dr TRAN, Viet (Institute of Informatics Slovak Academy of Sciences, Slovakia)

Presenter: Dr VERLATO, Marco (Istituto Nazionale di Fisica Nucleare - Sez. di Padova, Italy)

Session Classification: Supercomputing, High Throughput, Accelerator Technologies and Integration

Track Classification: Supercomputing, High Throughput, Accelerator Technologies and Integrations