

Comparison of Container-based Virtualization Tools for HPC Platforms.

Monday, September 23, 2019 5:45 PM (15 minutes)

Virtualization technologies are a fundamental element in cloud computing. Docker is the most known and used container platform worldwide. It is designed for microservices virtualization and application delivery but its model does not fit well with High-Performance Computing (HPC) platforms. HPC environments are multi-user systems where users should only have access to their own data and computing resources. Misconfigured Docker installations pave the way for privilege escalation, including the ability to access other users' data and, at the same time, gaining control of the cluster and computing resources.

In the world of HPC, the focus of containerised applications is not necessarily on DevOps, but on the ability to minimise HPC node configuration and manage applications' software dependencies through containers. Several open source initiatives have addressed this problem of bringing containers to the HPC space such as Singularity, Shifter, CharlieCloud and uDocker. In this sense, Singularity seems to be the most popular container system for HPC centres, but there are alternatives such as uDocker that support the execution of containers in user space, a key feature in HPC platforms. Therefore, it is important to analyze the benefits and drawbacks of these solutions when they are deployed in real HPC system and applied to scientific production applications.

All these tools, with potentially similar characteristics, bring the benefits of the containers to the HPC world. However, it is important to analyze important metrics in order to determine the advantages of one over another. The fields to analyze include, but are not limited to: interaction with Docker, support for Graphics Processing Unit (GPU), support for low-latency interconnects such as InfiniBand, support for Message Passing Interface (MPI), security and portability, privilege model, integration with Local Resource Management Systems (LRMS), among others. The objective of this communication is to show the behaviour and limitations of different container technologies in the context of HPC systems.

Keywords: virtualization, HPC, uDocker, Singularity, comparison, metrics.

¹ DevOps (Development and Operation) refers in this context to continuous integration/continuous delivery (CI/CD).

Primary authors: Ms NARANJO DELGADO , Diana María (UPV); MOLTÓ, Germán (Universitat Politècnica de València); GOMES, Jorge (LIP); DAVID, Mário (LIP); BLANQUER ESPERT, Ignacio (Universitat Politècnica de València)

Presenter: Ms NARANJO DELGADO , Diana María (UPV)

Session Classification: IBERGRID Contributions

Track Classification: R&D for computing services, networking, and data-driven science.