Contribution ID: **3** Type: **not specified**

Sarus: Highly Scalable Docker Containers for HPC Systems

Thursday, 20 May 2021 12:00 (30 minutes)

Sarus is a container engine for HPC systems that provides the means to instantiate high-performance containers from Docker images.It has been designed to address the unique requirements of HPC containers, such as integration with hardware accelerators, quick deployments at scale, security and permissions enforcement on multi-tenant hosts, and parallel filesystems. Sarus leverages the Open Container Initiative (OCI) specifications to extend the capabilities of a standard runtime through dedicated hook programs, implementing container customization at deployment time and without the user's intervention.

This presentation will highlight how OCI hooks can enable portable, infrastructure-agnostic images to achieve native performance on top of HPC-specific devices such as GPUs and high-speed interconnects.

Thanks to their standalone nature and standard interface, OCI hooks can be independently developed to target specific features, and can be configured according to the characteristics of particular host systems. The same container image can thus be used across the whole development workflow, from early tests on a personal workstation to deployments at scale on flagship systems, while benefiting from the advantages of each platform.

Presenter: MADONNA, Alberto (CSCS)