

Performance Characterization of Containerized HPC Workloads

Wednesday, 22 March 2023 14:40 (20 minutes)

Approaching the exascale era, complex challenges arise within the existing high performance computing (HPC) frameworks. Highly optimized, heterogeneous hardware systems on the one side, and HPC-unexperienced scientists with a continuously increasing demand for compute and data capacity on the other side. Bringing both together would enable a broad range of scientific domains to enhance the models, simulations and findings while efficiently using the existing and future compute capabilities. Those systems will continue to develop a more and more heterogeneous landscape of compute clusters, varying classical computation and accelerator cores, interconnects or memory and storage protocols and types. Adapting user applications to those changing characteristics is laborious and prevents enhancing the core functions of the applications while focusing on deployment and runtime issues. Consequently, containerization is one key concept to shift the focus back to the actual domain science, removing incompatible dependencies, unsupported subprograms or compilation challenges. Additionally, an optimized efficiency of the compute systems' usage is reachable, if system owners would be aware of the actual requirements of the containerized applications.

Our proposed work provides a methodology to determine, analyze and evaluate characteristic parameters of containerized HPC applications to fingerprint the overall performance of arbitrary containerized applications. The methodology comprises the performance parameter definition and selection, a comparison of suitable measurement methods to minimize overhead, and a fingerprinting algorithm to enable characteristics comparison and mapping between application and target system. By applying the methodology to benchmark and real-world applications we aim to demonstrate its capability to reproduce expected performance behavior and build prediction models of the application's resource usage within a certain trash-hold. We enable a twofold enhancement of today's HPC workflows, an increase of the system's usage efficiency and a runtime optimization of the application's container. The system's usage efficiency is enabled by container selection and placement optimizations based on the container fingerprint, while the runtime profits from a streamlined, target-cluster-oriented allocation and deployment to optimize the time-to-solution.

The adaption module and the most prospective technology to overcome endless adaption of the application's program code is containerization, which offers portability among heterogeneous clusters and unprecedented adaptability to target cluster specifications. Containers like Singularity, Podman or Docker are well known for cloud usage and micro-service environments. During the last years containers like Apptainer or Charliecloud became also widespread in certain HPC domains, since their capabilities to include high data throughput, intra- and inter-node communication, as well as the overall scalability increased enormously.

We base our approach on the EASEY (Enable exASclae for EverYone) framework, which can automatically deploy optimized container computations with negligible overhead. Today's containers are natively not able to automatically use all given hardware at best, since the encapsulated application varies on computing, memory or communication demands. An added abstraction layer, although enabling many programming models and languages to be executed on very different hardware, is not able to make use of all provided hardware features. An enhanced EASEY framework will support distinct optimization tunings without any human interaction during compilation of the container. Based on the introduced methodology we will demonstrate how these optimizations could impact the performance of the containerized HPC application.

Primary author: HÖB, Maximilian (Leibniz-Supercomputing Centre (LRZ))

Presenter: HÖB, Maximilian (Leibniz-Supercomputing Centre (LRZ))

Session Classification: Converging Infrastructure Clouds, Virtualisation & HPC

Track Classification: Track 9: Converging High Performance Computing Infrastructures: Supercomputers, clouds, accelerators