

On the Implementation of MPI Cluster as a Service on Supercomputer System

Wednesday, 10 October 2018 17:00 (15 minutes)

The vast majority of HPC users are heavily leveraging MPI middleware for their applications. Historically, MPI was mainly configured on Supercomputer Systems and the applications were living in the boundaries set by the system administrators. This led to different issues, including but not limited to problems with application distribution, environment configuration, resource allocation and filesystem permissions. Recently, the expansion of Cloud Computing brought the attention of many HPC users to offerings like Infrastructure as a Service, Platform as a Service, Software as a Service and HPC as a Service. These Services give the power users much more granular control over the provided resources.

For the last year we have been researching variety of Linux operating system-level virtualization technologies aiming to mimic the flexibility, isolation and resource management provided by Cloud Computing into world of Supercomputer Systems without compromising the performance.

Our research resulted in Linux containers which gained their popularity and were adopted due to their small footprint, distribution form, runtime isolation and relatively neglectable performance overhead. These make them a very good candidate for implementing virtual Supercomputer Systems.

In this talk we present the approach that we used to provide our users with the ease to deploy virtualized MPI clusters and the power to control their configurations through the associated lifecycle operations. We framed all these in MPI cluster as a Service solution.

For platform implementation we used Supercomputer System Avitohol at IICT-BAS which is the core of the scientific computing infrastructure in Bulgaria and currently the most powerful supercomputer in the region with its 150 computational servers each equipped with two Intel Xeon Phi coprocessors and theoretical peak performance of 412.3 TFlop/s in double precision.

Our experiments showed that the performance overhead of executing MPI applications inside MPI Linux container-based cluster is close to zero. Hardware capacity is used more effectively by many concurrent users. MPI programs can be developed and sanity tested on local computer and easily transferred to the Supercomputer Systems. By using Linux containers, we have improved the overall Quality of Service for Avitohol users of scientific computing.

The application domain of this design is not limited to HPC but IoT, meteorology, traffic control, trading systems, in other words almost any MPI application available today.

Type of abstract

Presentation

Summary

In this talk we present our results on implementation of MPI cluster as a service on a supercomputer system without compromising the performance. For platform implementation we used Supercomputer System Avitohol which is the core of the scientific computing infrastructure in Bulgaria. Our experiments showed that the performance overhead of executing MPI applications inside MPI Linux container-based cluster is close to zero. Hardware capacity is used more effectively by many concurrent users. The application domain of this design is not limited to HPC but IoT, meteorology, traffic control, trading systems, etc.

Primary author: Mr SIMCHEV, Teodor (IICT-BAS)

Co-authors: Prof. KARAIVANOVA, Aneta (IICT-BAS); Dr ATANASSOV, Emanoil (IICT-BAS); Dr GUROV, Todor (IICT-BAS)

Presenter: Mr SIMCHEV, Teodor (IICT-BAS)

Session Classification: Computing Services Part II

Track Classification: Area 3. Computing and Virtual Research Environments