

BigHPC: A Management Framework for Consolidated Big Data and High-Performance Computing

Tuesday, November 1, 2022 10:50 AM (25 minutes)

The BigHPC project is bringing together innovative solutions to improve the monitoring of heterogeneous HPC infrastructures and applications, the deployment of applications and the management of HPC computational and storage resources. It aims as well to alleviate the current storage performance bottleneck of HPC services.

The BigHPC project will address these challenges with a novel management framework, for Big Data and parallel computing workloads, that can be seamlessly integrated with existing HPC infrastructures and software stacks.

The BigHPC project has the following main goals:

- Improve the monitoring of heterogeneous HPC infrastructures and applications;
- Improve the deployment of applications and the management of HPC computational and storage resources;
- Alleviate the current storage performance bottleneck of HPC services.

BigHPC platform is composed by the following main components:

- Monitoring Framework
- Virtualization Manager
- Storage Manager

For the BigHPC project, the main mission of the Monitoring Framework component is to empower users with a better understanding of their jobs workload and to help system admins to predict possible malfunctions or misbehaved applications. BigHPC will provide a novel distributed monitoring software component, targeted for Big Data applications, that updates the state of the art of previous solutions, by:

- supporting Big Data specific metrics, namely disk and GPU; being non-intrusive, i.e., it will not require the re-implementation or re-design of current HPC cluster software;
- efficiently monitoring the resource usage of thousands of nodes without significant overhead in the deployed HPC nodes;
- being able to store long-term monitoring information for historical analysis purposes;
- providing real-time analysis and visualization about the cluster environment.

Virtual Manager (VM) is a component in the BigHPC implementation that aims to stage and execute application workloads optimally on one of a variety of HPC systems. It mainly consists of two subcomponents, ie. VM scheduler and VM repository.

The Virtual Manager Scheduler provides an interface to submit and monitor application workloads, coordinate the allocation of computing resources on the HPC systems, and optimally execute workloads by matching the workload resource requirements and QoS specified by the user with the available HPC clusters, partitions and QoS reported by the BigHPC Monitoring and Storage Manager components respectively.

Additionally, the Virtual Manager Repository provides a platform to construct and store the software services and applications that support BigHPC workloads as container images. It then provides those uploaded images in a programmatic way when a workload request is submitted to the Virtual Manager Scheduler for execution.

The storage performance has become a pressing concern in these infrastructures, due to high levels of performance variability and I/O contention generated by multiple applications executing concurrently. To address the previous challenge, storage resources are managed by following a design based on Software-Defined Storage (SDS) principles, namely through a control plane and data plane. With an architecture tailored for the requirements of data-centric applications running on modern HPC infrastructures, it is possible to improve I/O performance and manage I/O interference of HPC jobs with none to minor code changes to applications and HPC storage backends.

In order to keep all development tasks in a common path, some good practices are needed to get a shorter development life cycle and provide continuous delivery and deployment with software quality. All these BigHPC components are being tested on two different testbeds: development and preview. In the development testbed there is a workflow to test each platform component, where a pipeline allows to automate all required tasks related to software quality and validation. Afterwards, the components are tested in real infrastructure using the preview testbed, where the integration and performance tests take place.

The implementation of software quality and continuous deployment adopts a GitOPS set of practices that allow the delivery of infrastructure as code and application configurations using git repositories. In this work we are creating the git workflow being adopted for application development and the tools that we are joining together to answer the three components of GitOPS: infrastructure as code, merging changings together and deployment automation.

In this presentation, we will do a brief introduction of the BigHPC project, but focusing on the main challenges we found during this project, facing the goals of the project and the reality of HPC BigData environments concerning the integration tasks.

Desired slot length

15

Speaker release

Yes

Presentation will be held...

in the conference venue

Author: Mr BERNARDO, Samuel (LIP)

Co-authors: Dr RUHELA, Amit (TACC); Dr ANTUNES, Bruno (Wavecom); Dr CAZES, John (TACC); Dr GOMES, Jorge (LIP); Prof. PAULO, João (INESC TEC); Mr COSTA, Júlio (Wavecom); Ms MIRANDA, Mariana (INESC TEC); Mr VIANA, Miguel (LIP); Dr DAVID, Mário (LIP); Prof. CASTRO, Nuno (LIP); Dr MACEDO, Ricardo (INESC TEC); Mr HARRELL, Stephen (TACC); Prof. CHIDAMBARAM, Vijay (UT Austin)

Presenter: Mr BERNARDO, Samuel (LIP)

Session Classification: Computing and Batch Services

Track Classification: Computing & Batch Services