

BigHPC

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

Advanced Computing

Bruno Antunes, Júlio Costa (Wavecom); João Paulo, Ricardo Macedo, Mariana Miranda, Paula Rodrigues (INESC TEC); John Cazes, Stephen Harrell, Amit Ruhela (TACC & UT Austin); Samuel Bernardo, Miguel Viana (LIP)



Challenges

- **Efficient support of both** parallel computing **and** Big Data workloads
- **Active management of available** computational **and** storage HPC resources
- Transparent application access **to such resources**
- Performance isolation **and** fairness **across these resources**

These are addressed with a novel management framework for Big Data and parallel computing workloads

Goals

1. **Improve the monitoring of heterogeneous HPC infrastructures and applications**
2. **Improve the deployment of applications and the management of HPC computational resources**
3. **Improve storage performance and management for HPC services**
4. **Integrate the BigHPC components and establish a CI-CD pipeline** for future enhancements

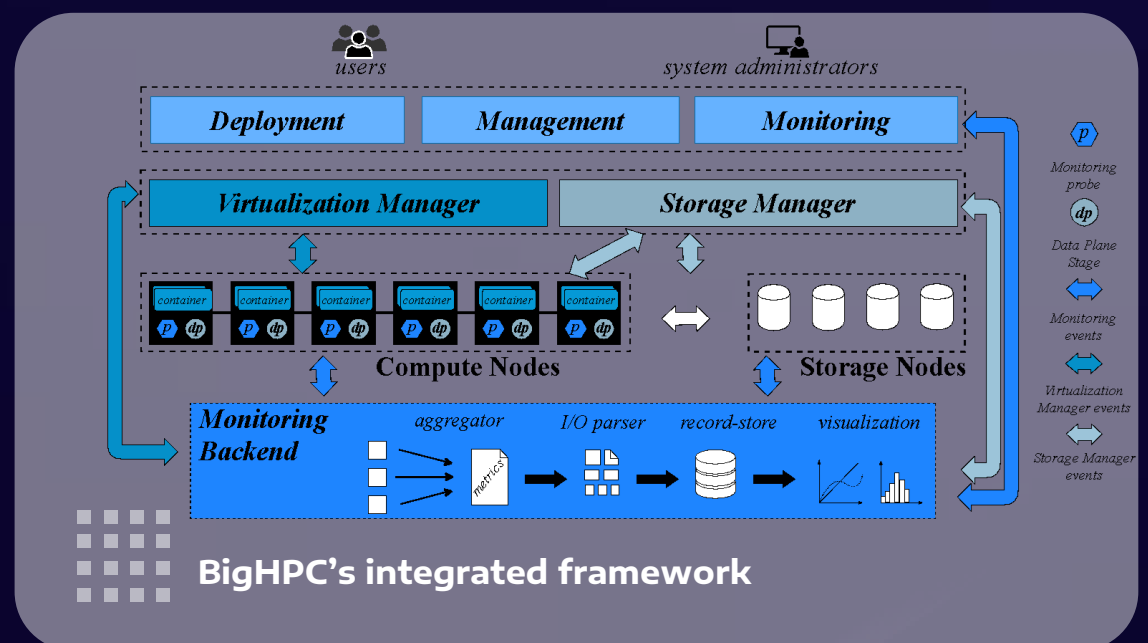
Final Architecture

The BigHPC's integrated framework is composed by three main components:

- A Virtualization Manager, responsible of storing/curating/managing the container files to deploy and also responsible to launching them in the cluster if there are available resources
- A Storage Manager, that will keep track of I/O requests being made to shared storage resources and will enforce QoS policies to improve I/O performance for applications using such resources.
- A Monitor Backend that will inform both the storage manager and virtual manager about the cluster status and workloads in real time, while storing data from past workloads for latter analysis

Prototypes

- **Virtualization Manager, 2023**
- **Monitor Backend**
 - HECTOR, 2021
 - MONICA, 2022
- **Storage Manager**
 - PAIO, 2021
 - Cheferd, 2023
 - PADLL, 2023



Selected Publications

- WineFS: a hugepage-aware file system for persistent memory that ages gracefully, SOSP 2021
- PAIO: General, Portable I/O Optimizations with Minor Application Modifications, Usenix FAST 2022
- TACCL: Guiding Collective Algorithm Synthesis using Communication Sketches, NSDI 23
- Taming Metadata-intensive HPC Jobs Through Dynamic, Application-agnostic QoS Control, CCGrid 2023



Fundings

