ALICE

Usage and perspectives of HPC in ALICE

Latchezar Betev, Irakli Chakaberia, Stefano Piano

CERN

BERKELEY LAB

INFN TRIESTE

General principles for HPC integration and use

- HPCs are treated as integral part of the ALICE Grid infrastructure
  - Subject to same rules as the other Grid sites
  - Made (to the extent possible) into a Grid node
  - ***All users and workflows are shielded from HPC specifics***
  - Manual support is not favoured - the Grid should be automatic
- Individual discussions with HPCs to reach a compromise on the infrastructure limitations
  - For example, network access and software delivery (more on that at the end)
  - Weigh the cost of integration with resources gain
  - Develop ‘custom solutions’ if warranted
  - Involve as much as possible interested third parties in the integration, i.e. students

Current utilization of HPCs in ALICE

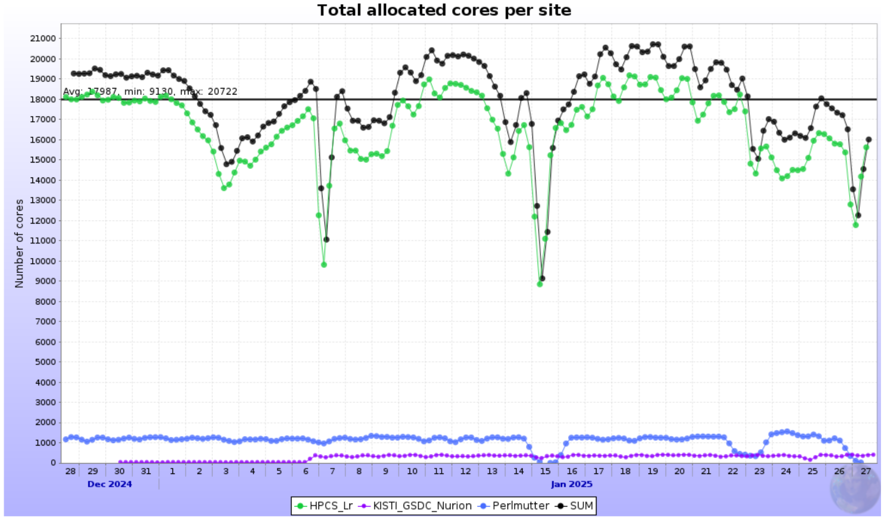

- 5% in average, 10% max contribution to total CPU allocation for ALICE

- 3 HPC successfully integrated through JAliEn in the ALICE Grid, following the principles from the previous slide
  - Lawrencium and Perlmutter at LBNL
  - Nurion at KISTI
- Co-located at a T2 and T1
- Profiting from direct connection with the local storage elements
- Able to run all types of ALICE workflow - reconstruction, simulation, analysis
- All deployed as ‘whole node’ - the payload management and resources allocation entirely managed by ALICE
  - Allows to respect and adjust to specific resources limits, for example network throughput

Some workload and development examples

- RAW data processing tests on Lawrencium and Nurion
  - Through remote data access (from CERN) on first, local data on second
  - A path to scale out a T1 resources and to run T0/T1 workflow on a T2
- Use of accelerators
  - Analysis with ML components running on Perlmutter
  - Ongoing development to increase GPU component in RAW reconstruction will increase the GPU resources use
  - Possibly greater contribution when MC transport becomes GPU-enabled
- Non-standard software delivery approach
  - Use of local fast storage to distribute software, eliminating need for CVMFS cache on Nurion

ALICE

HPC - a common approach

- High-Performance Computers pose unique integration challenges for WLCG
- By design, most of these are isolated, use proprietary systems and specialized software which clash with the Grid's interconnected, standardized environment
    - These unique characteristics require individual approaches and non-trivial amount of development and support
    - Today, their integration is up to the individual VOs efforts
- Key elements for a smoother HPC adoption
    - Common and standardized access mechanisms (similar to Grid gateways)
    - Standardized authorization and authentication tools and protocols
    - Sufficient external network capabilities for remote storage and services
    - Common methods for software delivery (e.g., CVMFS).
    - Use of modern techniques for software isolation and control (e.g., containers)
    - All of the above is easier to achieve if taken into account in the HPC design phase