

Managing Heterogeneous Compute Resources through Harvester and iDDS for ATLAS

Tuesday, 25 August 2020 16:30 (30 minutes)

The ATLAS experiment has been making significant efforts to integrate new resources such as HPC and pre-emptible cloud resources in addition to the traditional grid resources, to meet continuously increasing needs for data processing. It is difficult to optimally exploit all resources since their intrinsic nature and requirements are quite different.

Harvester has been developed around a modular structure to separate core functions and resource-specific plugins to simplify the operation with heterogeneous resources and provide a uniform monitoring view, enabling more intelligent workload management and dynamic resource provisioning based on detailed knowledge of resource capabilities and their real-time state. Harvester has been in production since early 2018 with various resources, which was one of the crucial milestones for LHC Run3.

Intelligent Data Delivery Service (iDDS) is an experiment agnostic service to orchestrate workload management system and data management, in order to transform and deliver data and let clients consume data just in time. iDDS has been actively developed by IRIS-HEP and ATLAS. One of the main goals of iDDS is to address performance issues and suboptimal resource usage in ATLAS workflows.

This talk will report architecture overview of Harvester and iDDS, migration of the entire ATLAS grid to Harvester, demonstration of scalable resource management with Kubernetes plugins, seamless integration of US and European HPC, and plans for the future.

Primary author: Mr LIN, Fa Hui (UTA)

Presenter: Mr LIN, Fa Hui (UTA)

Session Classification: Physics & Engineering

Track Classification: Physics (including HEP) and Engineering Applications