



Contribution ID: 271

Type: Oral

Large-scale HPC deployment of Scalable CyberInfrastructure for Artificial Intelligence and Likelihood Free Inference (SCAILFIN)

Monday, 4 November 2019 11:15 (15 minutes)

The NSF-funded Scalable CyberInfrastructure for Artificial Intelligence and Likelihood Free Inference (SCAILFIN) project aims to develop and deploy artificial intelligence (AI) and likelihood-free inference (LFI) techniques and software using scalable cyberinfrastructure (CI) built on top of existing CI elements. Specifically, the project has extended the CERN-based REANA framework, a cloud-based data analysis platform deployed on top of Kubernetes clusters that was originally designed to enable analysis reusability and reproducibility. REANA is capable of orchestrating extremely complicated multi-step workflows, and uses Kubernetes clusters both for scheduling and distributing container-based workloads across a cluster of available machines, as well as instantiating and monitoring the concrete workloads themselves.

This work describes the challenges and development efforts involved in extending REANA and the components that were developed in order to enable large scale deployment on High Performance Computing (HPC) resources.

Using the Virtual Clusters for Community Computation (VC3) infrastructure as a starting point, we implemented REANA to work with a number of differing workload managers, including both high performance and high throughput, while simultaneously removing REANA's dependence on Kubernetes support at the workers level. Performance results derived from running AI/LFI training workflows on a variety of large HPC sites will be presented.

Consider for promotion

No

Primary authors: HILDRETH, Mike (University of Notre Dame (US)); HURTADO ANAMPA, Kenyi Paolo (University of Notre Dame (US)); Mr KANKEL, Cody (University of Notre Dame); Mr HAMPTON, Scott (University of Notre Dame); Dr BRENNER, Paul (University of Notre Dame); Ms JOHNSON, Irena (University of Notre Dame); SIMKO, Tibor (CERN)

Presenter: HILDRETH, Mike (University of Notre Dame (US))

Session Classification: Track 9 –Exascale Science

Track Classification: Track 9 –Exascale Science