HIGH PERFORMANCE COMPUTING MODERNIZATION PROGRAM

RESEARCH PROJECT #: HPCMP-HIP-24-012

Bridging HPC, Kubernetes, and Cloud Environments for Scientific Simulation Workloads

About AFRL:

Air Force Research Laboratory (AFRL) is a scientific research organization operated by the United States Air Force Materiel Command. AFRL is dedicated to leading the discovery, development, and integration of aerospace warfighting technologies, planning, and executing the Air Force science and technology program, and providing warfighting capabilities to United States air, space, and cyberspace forces.

The composites performance team at the Air Force Research Laboratory Materials and Manufacturing Directorate uses a combination of novel and high-impact experiments, in-house high-fidelity HPC simulation software, and machine learning to characterize and predict the performance of current and emerging materials.

RESEARCH LOCATION: Wright-Patterson AFB, OH

PROJECT DESCRIPTION:

Exquisite supercomputers have been the target of high-fidelity scientific simulation for the past three decades, due to their extreme throughput, interconnect, and I/O. More recently, Kubernetes clusters are increasingly used for general workloads, including machine learning workloads. Finally, Google cloud provides scalable computational resources in a manner that is ideal for short-term or sparse projects. However, porting a job generally designed for classical HPC systems to Kubernetes or cloud requires significant effort. This project aims to define strategies for how to port specific types of scientific simulation jobs and develop simple tools that emulate the HPC environment familiar to many scientists.

The overall goal for the project is to determine techniques for effectively leveraging a variety of computing environments, capitalizing on massive recent investments at the lab and directorate levels. This will not replace the critical role of DoD HPC resources, but it will instead enable projects to leverage the most suitable resources available without the steep learning curve. The approach for the project focuses on 1) identification of the workflow and configuration details for effectively using DoD HPC, Kubernetes, and Google cloud environments for scientific simulation, including targeted parametric studies exploring the impact of configuration parameters, and 2) development of simple scripts that emulate the traditional DoD HPC tools across the other environments.

To achieve the goals of the project, the intern will:

Week 1: Complete in-processing, obtain access to a DoD HPC system, and learn how schedule simulations on an HPC system

Week 2: Learn about containerization and how to build custom Docker containers

Week 3: Learn how to submit jobs within a Kubernetes cluster using the Kueue or Volcano frameworks

Week 4: Learn how to create a virtual machine instance, configure the settings/network within the Google Compute Engine, and use a Google cloud VM to perform a scientific simulation

Week 5: Explore methods for transferring data between a client and each computing environment and transferring data across environments

Weeks 6-8: Develop scripts (within Python or another proficient language) to automate the process of porting scientific simulations between computing environments

Weeks 9-10: Author a report summarizing the work, document all code, and give a research presentation to research team and another to the broader AFRL RX community

In addition to the activities related to the project directly, the intern will have the opportunity to attend seminars focused on computing, machine learning, and material science; attend technical meetings across a variety of disciplines; participate in tours in the computing and material labs; and network with experts across disciplines. These activities will give the intern the opportunity to lead a research project typical to those in government labs, develop an understanding of a variety of computing environments, and show how tool development can impact the broader community of researchers in the lab. If desired, the intern can choose to pursue DoD technical report.

ANTICIPATED START DATE:

May 2024 – Exact start dates will be determined at the time of selection and in coordination with the selected candidate.

QUALIFICATIONS:

The ideal candidate should be:

- Enrolled in a computer science/engineering program or another STEM field coupled with a strong computing background
- Proficient in at least one well-known programming language (Python, C++, etc.)
- Be at least minimally familiar with Linux

Qualifications that are not required but would be helpful include:

- Familiarity with common Linux commands and shell scripting
- Familiarity with containers technology
- Familiarity with job scheduling and typical workflows in HPC environments
- Familiarity with Kubernetes and/or Google cloud.

ACADEMIC LEVEL:

Degree received within the last 60 months or currently pursuing:

- Bachelor's
- Master's
- Doctoral

DISCIPLINE NEEDED:

- Computer, Information, and Data Sciences
- Engineering
- Mathematics and Statistics
- Physics
- Science and Engineering related