Reference ID: 11227084408_Chan

Submission Date and Time: 12/16/2019 5:00:19 PM

Author Names & Affiliations

Submitting author: Chi-kwan Chan - Astronomy Department and Data Science Institute, University of Arizona

Additional authors: Parker Antin (Cellular and Molecular Medicine and CyVerse, UArizona); Tony Edgin (CyVerse, UArizona); Blake Joyce (Research Computing, UArizona); David LeBauer (Data Science for the Arizona Experiment Station, UArizona); Eric Lyons (School of Plant Scienc

Contact Email Address (for NSF use only): (hidden)

Research domain(s), discipline(s)/sub-discipline(s)

Cellular and Molecular Biology, Pharmacy, Agriculture, Plant Sciences, Astronomy, Physics, Data Science, and Research Computing/Cyberinfrastructure

Title of Response

Multi-Cloud Cyberinfrastructures and Software Stacks for Sustaining Data-Intensive Scientific Collaborations

Abstract

As data volume and complexity grow, traditional High-Performance Computing, cloud computing cyberinfrastructure, and software packages alone cannot sustain scientific innovation and discoveries. Setting up project-specific multi-cloud CI to address the specific needs of a project is now common for large collaborations. Our team identifies that there is a pressing need to invest research and development in i) Cloud-HPC convergence, ii) multi-cloud CIs, and iii) software stacks that integrate workflows and data management on multi-cloud CIs for data-intensive science projects. We suggest investing in hardware accelerators such as GPUs, FPGAs, and fast solid state drive (SSD) storages, to make cloud CIs more HPC-like; container orchestration frameworks for HPC; developing Kubernetes-based multi-cloud provisioning tools that adopts industrial best practices; systematic studies on using container technology for reproducibility in science; developing profiling services to monitor software performance and efficiently distribute them in the multi-cloud CIs; developing a fully integrated software stack to tackle the data handling problems; project management tools and 24/7 services; and inter-field software discovery.

Question 1 (maximum 400 words) – Data-Intensive Research Question(s) and Challenge(s). Describe current or emerging data-intensive/data-driven S&E research challenge(s), providing context in terms of recent research activities and standing questions in the field. NSF is particularly interested in cross-disciplinary challenges that will drive requirements for cross-disciplinary and disciplinary-agnostic data-related CI.

As data volume and complexity grow, traditional High-Performance Computing (HPC), cloud computing cyberinfrastructures (CIs), and software packages alone cannot sustain scientific innovation and discoveries. Modern data-intensive research demands many algorithms to work in unison on a wide range of hardware and across different data centers. Investments needs to flow into the following unmet needs in today's research and development: i) Cloud-HPC convergence, ii) multi-cloud CIs, and iii) integration software stacks for data-intensive workflows and management across multi-cloud CIs. Take the Event Horizon Telescope (EHT) collaboration for example. The international collaboration known best for publishing the first images of a black hole, offering direct evidence for the existence of an event horizon, and testing Einstein's general theory of relativity in strong gravity regime, was struggling with the problems common to data-intensive research today. The 5 petabytes of data from 8 geographically dispersed telescopes required data to be couriered to correlation facilities in Boston and Bonn, Germany. Multiple versions of the 5 terabytes correlated data sets were uploaded to CyVerse and Google Cloud for the distributed EHT team to work on, creating more data products. To alleviate the data intensive problems, the EHT capitalized on the scale separation in data volume along its pipeline to utilize different data centers. The EHT's in-house CyVerse--Google multi-cloud CI utilizes HPC and Cloud, brings computation in close proximity to the data, speeds up data processing, and enables the collaboratively process of the data and imaging black holes. Project-specific multi-cloud work-around CI like the ones described in the EHT above to address specific needs of a project is common in large collaborations, such as the TERRA Phenotyping Reference Platform in agriculture, and the Laser

Interferometer Gravitational-Wave Observatory (LIGO) and Large Synoptic Survey Telescope (LSST) projects in astronomy. In fact, the NSF funded CyVerse evolved from the domain-specific iPlant Collaboration. While it is agreed that the most innovative CI and software ideas come “near cutting-edge science”, it is also agreed that project-driven CI developments are difficult to get funded, thereby taking up the very limited personal time and resources of the scientists and distracts from scientific work. Thus, there is a consensus across fields of the difficulty in i) utilizing both HPC and Cloud, ii) deploying a multi-cloud CI to fulfill the data and compute needs, and iii) running complex workflows and manage distributed data products on a multi-cloud CI.

Question 2 (maximum 600 words) – Data-Oriented CI Needed to Address the Research Question(s) and Challenge(s). Considering the end-to-end scientific data-to-discovery (workflow) challenges, describe any limitations or absence of existing data-related CI capabilities and services, and/or specific technical and capacity advancements needed in data-related and other CI (e.g., advanced computing, data services, software infrastructure, applications, networking, cybersecurity) that must be addressed to accomplish the research question(s) and challenge(s) identified in Question 1. If possible, please also consider the required end-to-end structural, functional and performance characteristics for such CI services and capabilities. For instance, how can they respond to high levels of data heterogeneity, data integration and interoperability? To what degree can/should they be cross-disciplinary and domain-agnostic? What is required to promote ease of data discovery, publishing and access and delivery?

In this RFI, we suggest to bridge the gap between specific science projects and generic CI development by investing in i) Cloud-HPC convergence, ii) multi-cloud CIs, and iii) integration software stacks for data-intensive workflows and management. These investments will create common hardware and software resources for scientific projects, and passing on successful CI from large collaborations to smaller teams towards a blossoming ecosystem of holistic multi-cloud CIs. These investments will directly address the problems mentioned in question 1, and ultimately optimize the usage of NSF-funded computing resources, support data-intensive science projects, allow scientists to focus on the science, speed up data-intensive workflows, and sustain science innovations and discoveries. There are two aspects of i) Cloud-HPC convergence. First, on the CI/hardware front, we suggest investing in hardware accelerators such as GPUs and FPGAs, as well as fast solid state drive (SSD) storages, to make cloud CIs more HPC-like. We suggest deploying these hardware on both project-specific testbeds and to existing well-established CI such as CyVerse and Jetstream. The goal here is to bring the hardware diversity of NSF-funded CIs to the level of commercial Clouds. By making more hardware accelerators and faster storage options available on NSF-funded resource, researchers will be able to better utilize NSF-funded resources. Second, on the software front, we suggest investing in container orchestration framework for HPC. Traditional HPC supports tightly-couple computing jobs, while Cloud supports various data and compute needs of projects. Container orchestration framework for HPC will help Cloud-HPC convergence, and ultimately allows for tighter integration of HPC resources into multi-cloud CIs (see also next point). For ii) multi-cloud CIs, we suggest developing high-level multi-cloud provisioning tools that adopts best practices from the industry including infrastructure-as-code, gitops, and multi-cloud federation. Our vision of future data-focused CI is that, by editing a few YAML files, any team can launch a project-specific multi-cloud CI that fulfill the project's needs. Currently, projects such as the EHT require dedicated personnel, sometimes not paid for doing system admin, to maintain a multi-cloud CI. By funding the development of multi-cloud provisioning tools, we consolidate CI maintenance into common software projects, create better data-focused CIs, and free scinsitsits from doing system admin works. We further suggest building the multi-cloud software stacks on top of Kubernetes, a production-grade container orchestration framework developed originally by Google and widely supported by the industry. This offloads the low-level develop to the industry and enables us to focus on the science-driven aspects. There are three levels of research and development needed for iii) software stacks that integrate workflows and data management on multi-cloud CIs. Although controversial, the term "reproducibility crisis" is well known in the fields of psychology, medicine, and recently data science. Container technologies have the potential to address this. Nevertheless, most container best practices are developed by the tech industry and do not address the specific needs of science. Therefore, at the lowest level, we suggest investing on systematic studies to determine how to best use containers technology such as Docker and Singularity for reproducibility in science. At the middle level, we suggest developing online profiling tools to monitor runtime performance characteristics of scientific software on different hardware. The profiling results can be used to efficiently distribute workloads and data across different machines in a multi-cloud CI setup. With container technologies, it is possible to serve these profiling tools as standard services of the CI. At the highest level, we suggest developing a fully integrated software stack to tackle the data handling problems by providing exceptional and scalable capabilities for data transformation, version control, management, and reproducibility.

Question 3 (maximum 300 words) – Other considerations. Please discuss any other relevant aspects, such as organization, processes, learning and workforce development, access and sustainability, that need to be addressed; or any other issues more generally that NSF should consider.

In addition to our responses in Question 2, we have identified the following two underinvested areas: i) project management and 24/7 services; and ii) inter-field software discovery. For point i), While NSF-funded Cloud computing CIs such as CyVerse and Jetstream are designed for scientific workloads such as data analysis; however, they are not designed to serve 24/7 services such as single-sign-on, source code repositories, issue tracking system, or project management tools. Currently, small research teams often use free services such as ORCiD and GitHub and large collaborations often run their own servers or use commercial clouds. This creates overheads for research teams and again distracts scientists from working on the science. We plan to integrate these project management tools and services into the proposed multi-cloud provisioning tools. In addition, we suggest broadening the use cases of NSF-funded CIs to support scinece project management. For point ii), we suggest supporting a container registry for science, with machine- and human-readable metadata to help scientists discovery research software across fields. This is motivated by the FAIR (findable, accessible, interoperable and reusable) data principles. While software dependence made software distribution difficult, with container technology, it is now possible to make software FAIR and work "out of the container". The proposed container registry may live on existing NSF-funded CIs such as CyVerse and integrate with existing software efforts such as the Astrophysics Source Code Library.

-- End Submission –