Federal Register Notice 86 FR 46278, https://www.federalregister.gov/documents/2021/08/18/2021-17737/request-for-information-rfi-on-an-implementation-plan-for-a-national-artificial-intelligence, October 1, 2021.

Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource: Responses

Response to “Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource” (document 86 FR 39081)

National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign has been developing AI-capable resources for its own and external research for many years. Based on this experience, in this response we are providing information in response to Questions 1 and 4 jointly, in two aspects that are both related to roadmap elements D and H. Specifically, we have learned that for both technical (roadmap aspect D) and financial reasons (roadmap element H), an Artificial Intelligence Research Resource, whether providing this capability to a single institution such as the University of Illinois or to the Nation:

1. *It is essential to make use of multiple technology types and usage models of resources, including on-premises systems, national HPC and storage resources, and commercially-provided cloud resources, and as much as possible, to offer common interfaces to the resources.*

2. *It is equally essential to provide expertise, including technical support, alongside the computing and storage resources.*

In the remainder of this response, we provide more details on our experiences and explain how we can came to these conclusions.

Relevant NCSA experiences

1. Cloud investigations

During the last year NCSA has worked closely with many Fortune 500 companies to understand why and how some of them were moving to commercial cloud resources for their research activities, and to better understand the differences between on premises systems and cloud systems, both of which can be used for multiple purposes including AI research. During that time NCSA has run multiple proof-of-concept tests to find ways to move data science, machine learning, and normal high-performance computing solutions into the cloud. We took existing workflows and code and shifted them to Google Cloud Platform (GCP) to evaluate both performance and cost for these corporate partners. During this year, all basic functions of these partners' research workflows were moved and tested. We learned three key lessons:

1. The cloud platform was not built to run applications originally designed for high-performance resources. Much of the software used on cloud HPC environments is the same code developed for traditional HPC systems. However, support for that software within the cloud is typically not offered as part of the cloud solution. It is typically up to the individual to make connections to the developer of the code and to the community that supports it to run in the cloud. With the in-house solutions, the in-house staff has usually done the installation and can provide basic support for the software.
2. When testing performant file systems, access to knowledgeable storage engineers is more important when running in GCP than when using local hardware solutions.
3. When moving workflows into the cloud, we found shortcomings related to the debugging and error reporting needs of developing a new system. Running controlled benchmarks led to results that didn't match what was expected and debugging those solutions was more complex and close to impossible without planning for it in advance. Cloud resources do not keep logs and related information from one session to the next, which makes comparisons more difficult.

Overall, while the advantage of the unlimited resources seemed great on paper, the knowledge and skills to implement valuable science required much more technical staff time than advertised. On the other hand, using in-house platforms to obtain research results was much easier for graduate students and researchers without all the extra technical skills required.

Additionally, a small team at NCSA experimented with cloud computing in the context of the 2.4 PB TerraFusion project (https://digirolamo.web.illinois.edu/projects/terra-fusion/) in 2018, by re-running a subset of the project's processing and storing the data in AWS, then scaling the results from the subset to the entire project dataset. We learned that for this project, storage costs dominated our AWS expenses, and the cloud costs substantially exceeded the costs of equivalent services provisioned in-house.

2. HAL: an NSF-funded MRI resource for AI workloads

NCSA received an NSF Major Research Instrumentation (MRI) award to develop, deploy, and operate a purpose-built computational instrument for running AI workloads. The system, named HAL, has been operational since March of 2019, enabling over 700 users to develop and train AI models. The system consists of 16 IBM Power9 AC922 servers with NVIDIA V100 GPUs, Infiniband EDR interconnect and a DDN flash array storage. Coupled with the commercial and open-source software, the system enables researchers to achieve state-of-the-art results on many AI models. While developing and operating HAL, we learned a number of valuable lessons both about technology requirements for AI applications and challenges in using this technology to obtain state-of-the-art results.

1. The compute, storage, and interconnect components all have to be well-balanced and matched to enable a high-productivity environment. For example, our initial choice for the storage solution was only able to supply data fast enough to a handful of nodes, making it impossible to productively use the entire system. Only after deploying a substantially improved storage subsystem were we able to reap benefits of the entire system.
2. The AI software stack evolves at a very high rate, with new versions of major tools being updated literally every day. As a result, significant personnel effort is needed to keep the system current with the AI tools while maintaining compatibility and performance for developers and users.
3. The AI user community is very diverse, ranging from those developing domain-specific models to those developing actual AI systems. Our experience shows that the community needs a significant amount of help to utilize the technology efficiently. This help needs to come in a variety of forms, ranging from providing user-friendly interfaces to the complex underlying hardware and software, to old-fashioned one-on-one tutoring, troubleshooting, and software tool support. This level of user support is critical to enable the use of the AI technology.

3. Extreme-scale data analysis on Blue Waters

At-scale computing of data intensive workloads such as large-scale image analysis and AI inferencing requires support for a set of elements that don't appear at the scale of desktop, research group, departmental or campus scale systems. Projects such as Digital Elevation Model (DEM) generation and tree mapping in satellite imagery projects described below are examples where NCSA’s expertise in data movement, workflow software, software containerization, and HPC programming techniques enabled domain experts to achieve unprecedented results.

The DEM project (https://bit.ly/NCSA-AI-DEM), a collaboration of NCSA at the University of Illinois, the Polar Geospatial Center (PGC/UMN) at the University of Minnesota, the Ohio Supercomputer Center at the Ohio State University (OSC/OSU), and the National Geospatial-Intelligence Agency, generated millions of individual stereoscopic DEMs extracted from pairs of sub-meter resolution satellite imagery by applying fully automated, stereo auto-correlation techniques to overlapping pairs of high-resolution optical satellite images using the open source Surface Extraction from Triangular Irregular Network (TIN)-based Search-space Minimization (SETSM) software developed by OSC/OSU. Images were processed with SETSM using compute power on the NCSA Blue Waters supercomputer system using open-source Swift and Parsl workflow software for task management and sub-scheduling. Globus and AWS services were used to transfer petabytes of satellite imagery and resulting DEMs. Memory use and computation were optimized with assistance from staff HPC expertise and associated tools, all of which are techniques that also need be used in analyzing large-scale datasets for AI applications.

The Sub-Sahara tree mapping project (https://bit.ly/NCSA-AI-tree-mapping), a collaboration of Illinois/NCSA, PGC/UMN, and NASA's Goddard Space Flight Center, used machine learning algorithms and sub-meter satellite imagery to identify and measure the crown diameter of more than 1.8 billion trees across an area of more than 500,000 square miles, or 1,300,000 square kilometers. The success of the project relied heavily on NCSA's expertise with HPC software containerization and IO best-practices to exploit the power of a large-scale file system for efficient inferencing.

4. XSEDE user support and initial Delta user support

The Extreme Science and Engineering Discovery Environment (XSEDE, https://www.xsede.org) is an NSF-funded virtual organization that integrates and coordinates the sharing of advanced digital services, including supercomputers and high-end visualization and data analysis resources, with researchers nationally to support science. NCSA leads the XSEDE project, which is operated in collaboration with 16 other institutions[1]. There is a rapidly growing number of AI-based or AI-enhanced applications already being supported on XSEDE-allocated resources.

For the upcoming NSF-funded and NCSA-hosted Delta supercomputer system (http://www.ncsa.illinois.edu/enabling/delta), interviews with investigators who plan to use the system show movement in the science community from purely physics-based models to adding AI components for parts of the models or the analysis of the data from

[1] https://confluence.xsede.org/display/XT/Subaward+PIs

the models. These researchers typically don't have AI expertise or experience, but have decided that they need to learn these skills and use these methods to do the best research.

In both cases (XSEDE and Delta), these researchers need support, and our XSEDE experience shows that this support needs to be long-term (multiple years in many cases), in part because these researchers' needs change as they progress and learn more about their science and about potential AI methods.

Additionally, our XSEDE experience has shown that different researchers request different types of hardware resources, and that these needs change over time. Today, this typically focuses on GPUs, and specifically, many GPUs per node and large memory per node and per GPU. This leads to some XSEDE systems being in high demand for AI research, while other systems are more used for other types of computational and data research.

5. C3.ai Digital Transformation Institute

The C3.ai Digital Transformation Institute (C3.ai DTI, https://c3dti.ai/), a research consortium that includes a support effort conducted jointly by NCSA and UC Berkeley, has accumulated specific experience supporting research on enterprise-class cloud AI infrastructure. The C3.ai DTI is entering its second award year. From a user support perspective, a few key lessons learned after the first award year going into the second are:

- The ability to leverage multiple resources, including cloud, HPC and in this case, C3.ai, is crucial for broad research team success.
- In-house technical support that can work closely with research teams and steer them toward the appropriate platforms removes barriers for those teams.
- Close cooperation with C3 (or other cloud framework providers) is needed both to leverage their in-house expertise and, critically, to translate research goals into a company culture more accustomed to considering business goals and deliverables.
- Support teams need diverse expertise, including data science/AI/machine learning as well as software engineering, computer science, HPC, system admin and solid scientific backgrounds, to be able to provide successful support.

For C3 deployments, the close working relationship between the C3.ai DTI support team and C3.ai engineers has enabled us to build technical proficiency, develop working code, and deploy infrastructure for new research teams as they come online.

The software engineering, cloud and system administration, and diverse domain expertise of our support team have all been critical in enabling the C3.ai DTI to engage research groups on a scientific and computational level and to develop tools to facilitate quicker onboarding and prototyping on sophisticated cloud deployments. Any effort to support large scale AI research efforts will benefit from both diverse computational environments as well as a deep and broad support team expertise.

NCSA lessons learned

From these experiences, we have learned the following two lessons:

1. It is essential to make use of multiple technology types and usage models of resources, including on-premises systems, national HPC and storage resources, and commercially-provided cloud resources, and as much as possible, to offer common interfaces to the resources.

Over multiple decades, an diverse ecosystem of resources has become available to researchers, including on-premises systems (including company, university, and laboratory compute and storage resources, whether operated as clusters or clouds), national HPC and storage resources (e.g., provided by NSF, DOE), and commercially-provided cloud resources.

Each of these types of resources has positive and negative aspects for particular use cases, including the means of access and scheduling, the scale of usage, the ease of use, models of service provisioning (IaaS, FaaS, SaaS, etc.), specific resources (number and type of CPUs, GPUs, FPGAs, TPUs, etc.), data access, the cost to the end user, and the costs and business model for the resource provider, including support staff, system costs, operations, and energy usage.

While any one use case might have a resource that is best according to some criteria, there will never be a single best resource (or even model of resources) for all use cases. Therefore, a general National AI Research Resource should support and offer access to multiple models of resources, which includes resources that are now and likely will continue to be physically distributed. Similarly, data is now and will continue to be physically distributed, including both public and private data, and use cases that are data-intensive will likely be better served by computing where the data is than forcing a user to move the data to another resource. Finally, it is imperative to compare full costs when making choices about resource investment, and specifically, not to assume that commercial clouds will always be less expensive than other options, especially when coupled with the support expertise necessary to effectively be utilized.

This distributed set of resources will need to change over time in response to the evolution of both hardware and software for AI research, and elements of it must be able to support a variety of different types of work, including basic research in AI methods as well as applied research in the use of models to address both academic and commercial problems. This requires resources that include stable “production-AI” resources that can be used for method development and applied research as well as those that support the investigation of emerging novel technologies, including experimental platforms.

2. It is equally essential to provide expertise, including technical support, alongside the computing and storage resources.

There are a relatively small number of highly experienced AI users, while many potential users of AI have little understanding of AI beyond its promise, and some college graduates who have taken AI-related coursework are quite knowledgeable about AI methods, but typically have limited exposure and experience with AI tools, especially for any practical-scale applications.

These different categories of users have different needs, such as operational support for highly-experienced users, basic training and education in AI and the use of AI platforms for inexperienced users, and tool-specific training for those in the middle, particularly as tools change. These tools are frequently deployed across many platforms (edge-to-cloud-to-user), they consist of complex workflows, and are rapidly evolving. Overall, developing applications in such environments is very demanding and challenging. Therefore it is critical to have access to the technical expertise provided by dedicated personnel in order to best utilize the systems. Additionally, experienced support staff with AI expertise can help users of all types avoid common errors that could make their results incorrect or of limited utility (e.g., biased.)

In more detail, AI researchers who want to use large-scale shared resources are confronted with a myriad of options and possibilities. Depending on the goals of the researcher and the current state of their software development for research, cloud and other shared platforms can present both obstacles and opportunities. One common scenario is a team that has developed prototype software that runs on a single local computer (laptop or small workstation) and now wants to move it to a cloud. This team will typically need support in identifying, provisioning, and then deploying their software on the appropriate resources in that environment, which is typically not available from a cloud vendor but is for in-house resources and nationally-provided research HPC systems.

Support for research users is not the same as the general support that system and software vendors typically provide. Researchers work in an iterative, question-based manner, where they typically cannot provide requirements for what they need in advance, but rather, their requirements may change after each iteration of the research cycle. Therefore, the staff who support them need to be comfortable working in this same style. Over the last ten-plus years, new roles for staff have been developed who can do this, mixing understanding of researchers and their style with professional skills and understanding of underlying systems. These roles, such as research programming or research software engineering (RSE) that specialize in software development aspects, and data science that specializes in data analysis aspects, have become essential parts of the research landscape in academia, national laboratories, and industry, and have come together in the US Research Software Engineer Association (US-RSE, https://us-rse.org) and the Academic Data Science Alliance (ADSA, https://academicdatascience.org).

More complicated scenarios involve teams that lack the expertise and time to become experts in web platform management, database administration, and web development, all of which can comprise a comprehensive research platform and may become relevant as teams scale up their efforts. For smaller research teams this is untenable. While porting existing workflows onto large-scale shared resources can enable more throughput and immediate return for researchers, many teams can benefit from a more comprehensive approach in managing their research software projects without becoming experts in all of the underlying technologies.

The use of frameworks that integrate database, web, and computation can enable better end-to-end computational and analysis workflows, especially for large scale HPC-class problems, but also introduce more abstraction and somewhat steeper learning curves than basic porting of existing workflows. A curated approach for AI in large-scale shared resources that leverages a framework-like approach should offer many on-ramps for small to medium teams of domain-experts. Some will be content with a basic computation-only environment, but many teams will benefit from a more comprehensive approach that includes integrated analysis, hyperparameter tuning, cataloging of results, and so on.