

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

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government nor those of the National AI Research Resource Task Force., and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality, or content of all external links included in this document.



September 29, 2021

The National Science Foundation
Attn: Wendy Wigen, NCO
2415 Eisenhower Avenue
Alexandria, VA 22314
[REDACTED]

Re: RFI Response – National AI Research Resource (86 FR 39081)

Dear Ms. Wigen:

The University Corporation for Advanced Internet Development (d/b/a “Internet2”) is pleased to submit these comments in response to the National Science Foundation’s (“NSF”) request, in coordination with the Office of Science and Technology Policy (“OSTP”), for community input to inform the work of the National Artificial Intelligence Research Resource (“NAIRR”) Task Force.

BACKGROUND ON INTERNET2

Internet2 is a non-profit, member-driven advanced technology community founded in 1996 by the nation’s leading higher education institutions that provides a secure high-speed network, cloud solutions, research support, and identity and access management services tailored for research and education (“R&E”). Internet2 helps U.S. R&E organizations to solve shared technology challenges and develop innovative solutions in support of their educational, research, and community service missions. Internet2 also operates the nation’s largest and fastest coast-to-coast national research and education network (“NREN”), which now serves 323 U.S. universities, 59 government agencies, and 45 regional and state education networks. In addition, Internet2 collaborates with numerous leading corporations that work with the R&E community, as well as a multitude of NREN partners across the globe that represent more than 100 countries.

In addition, Internet2 operates the InCommon Federation, which facilitates secure, unified, and seamless single sign-on access to local and global research and academic collaboration resources for more than 10 million users and 800 educational institutions, research organizations, and commercial resource providers in the U.S.

InCommon makes possible trustworthy academic collaboration that reaches far beyond what a single organization can do on its own, through identity and access management technologies and services that are integrated across the globe.

Internet2 also offers eduroam service in the U.S to enable seamless roaming Wi-Fi at nearly 1,000 colleges, universities, schools, museums, libraries, and research facilities across the country.

Further, Internet2 plays a key role as a convener and facilitator of the R&E community. Internet2 regularly brings together representatives from academia, federal agencies, and private industry to foster collaboration and find solutions to common challenges related to R&E cyberinfrastructure that no single institution or organization could accomplish independently.

Finally, Internet2 supports the R&E community through the Internet2 NET+ Cloud Services Program, which enables R&E institutions in adopting cloud solutions through a streamlined process that minimizes the business, legal, financial, and other challenges associated with migrating from on-campus to cloud-based solutions.

INTERNET2'S RESPONSES TO SPECIFIC RFI QUESTIONS

1. What options should the Task Force consider for any of the roadmap elements A through I and why?

D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advanced computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.

To facilitate access to the envisioned advanced computing resources for researchers across the country, several key characteristics should be included in the establishment of the NAIRR:

- The NAIRR should support hybrid models for advanced computing resources to enable the broadest possible set technology advancements, scalability, and use cases ranging from batch processing to low latency real-time applications. The hybrid models should include on-premise, national centers, commercial clouds, in-network, and edge computing scenarios.
- Given that artificial intelligence (“AI”) and machine learning (“ML”) often require the use of large data sets, high-speed networking and data transfer

capabilities will be critical for all researchers across the country no matter their institution. Therefore, the NAIRR should support future federal efforts focused on funding such network capabilities across all higher education institutions.

- To enable global trustworthy academic AI collaborations, secure access should be based on the well-established model of Multilateral Federated Identity Management.
- To develop resident expertise, the NAIRR should include support for a variety of training programs and communities of practice for researchers, campus research computing and data professionals, and campus identity management professionals. Training the campus professionals is necessary to ensure a sustainable ecosystem of resident expertise for researchers.

E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the National Artificial Intelligence Research Resource.

Given hybrid models for the NAIRR, a fundamental barrier to the dissemination and use of high-quality government data sets is the lack of high-speed networking in many of the country's higher education institutions, especially those that have been chronically underserved. Consideration should be given to funding the necessary network infrastructure to connect all of these institutions. Additionally, commercial clouds are likely to be an important part of the NAIRR for both compute and storage resources. The associated data egress fees have been seen as a barrier by the research community. National and regional R&E networks have a long-proven track record of providing the networking infrastructure for data intensive research (including data egress waivers and private/direct connections for commercial cloud usage). Therefore, it is recommended that consideration be given to funding for the expansion and/or upgrading of R&E networks to ensure that all higher education institutions have adequate connectivity to the NAIRR resources and data sets.

F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its management of access controls.

The InCommon Federation (<https://incommon.org/>) implements security requirements for access controls based on several pertinent international standards for R&E federations. These should be considered for NAIRR access controls:

- [REFEDS MFA Profile](#), a means for RP/SP to request MFA and for the CSP/IdP to respond.

- [REFEDS Assurance Framework](#), especially for its identity proofing and credential issuance claims.
- [REFEDS R&S Entity Category](#), a program for release of basic user attributes in a data minimizing manner.
- [Security Incident Response Trust Framework for Federated Identity \(Sirtfi\)](#), a framework for collaborating in managing security incidents across R&E federations.

4. What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Secure Access Control:

The previously mentioned InCommon Federation is a well-established building block to support the NAIRR's needs for *secure access control*. InCommon facilitates secure, unified, and seamless single sign-on access to local and global research and academic collaboration resources for more than 10 million users and 800 educational institutions, research organizations, and commercial resource providers in the U.S.

InCommon makes possible trustworthy academic collaboration that reaches far beyond what a single organization can do on its own, through identity and access management technologies and services that are integrated across the globe.

Dissemination and use of high-quality government data sets:

To provide necessary high-speed network access to higher education and other research institutions and support the dissemination of high-quality government data sets, the NAIRR should leverage the ecosystem of state, regional, and national R&E networks. As the U.S. NREN, Internet2 provides private direct connections to commercial cloud providers and data egress waivers that can be used in support of the NAIRR. Additionally, Internet2 has collaborated with the NSF-funded Open Science Grid ("OSG") to host caching servers in the Internet2 national network, which, if expanded, could be used to disseminate high-quality government data sets.

Resident Expertise:

Internet2 provides various opportunities for training and project execution for identity management professionals at universities and research institutions, including [InCommon Academy](#) and the [InCommon Catalysts](#) program.

To address challenges and facilitate the most effective use of cloud technologies and platforms, Internet2 also offers the [Cloud Learning and Skills Sessions \(CLASS\)](#) program. CLASS assists U.S. higher education institutions, research IT organizations, and research groups by providing training to effectively leverage cloud technologies and platforms for research workflows. A combination of vendor-neutral guidance across cloud providers, and training on the tools and technologies supported by public cloud providers, allows a broad range of research use cases to utilize these important resources more effectively. Participants join the CLASS community of practice where they can share information and lessons learned.

5. What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public-private partnerships should not only play a role in the implementation of the NAIRR compute and storage resources, but they should play a role across all the capabilities required to implement the NAIRR.

Not-for-profit companies, such as R&E networks, have a long history of successfully partnering with government, academia, and the for-profit private sector to provide research cyberinfrastructure.

Examples can be found in NSF programs supporting cyberinfrastructure. The NSF Campus Cyberinfrastructure (CC*) program solicitation (NSF 21-528) Program-wide Criteria states:

The plan should include the campus status and plans with respect to *federated identity and specifically InCommon*, including: if the campus is registered with InCommon as supporting the Research and Scholarship (R&S) Entity Category to streamline integration with research applications.

NSF 21-528 also states under the Campus Computing and the Computing Continuum Program Area High-Performance Network Connectivity and Specification:

Proposals must describe the network connectivity of the proposed computing resource, both intra-campus [for example, the campus network path(s) connecting

the resource with the researchers and driving science applications on campus], and inter-campus (for example, showing the network path *connecting with the regional exchange point or Internet2*).

Furthermore, the NSF International Research and Education Network Connections (IRNC) program solicitation (NSF 20-535) states under the Infrastructure Improvement and Support Program Area – Open Exchange Points that:

Proposals must demonstrate a commitment to operation of an open exchange point, for example: support for homing of multiple international links; high capacity connectivity to Internet2, ESNNet, and, if relevant, FABRIC; providing maximum flexibility in connectivity and peering; automated, dynamic switching network services; and, in the best interests of the end users – the researchers, educators, and students in the U.S. – a demonstrated commitment to a productive partnership and collaboration with *Internet2, the primary NREN for the NSF community*.

Another example of public-private partnerships between Internet2, government, academia, and the for-profit private sector is the Cooperative Agreement (Award #1904444) between NSF and Internet2 called “Exploring Clouds for Acceleration of Science (E-CAS).” The Abstract states the following:

Internet2 leads the “Exploring Clouds for Acceleration of Science (E-CAS)” project in partnership with representative commercial cloud providers to accelerate scientific discoveries. The effort seeks to demonstrate the effectiveness of commercial cloud platforms and services in supporting applications that are critical to growing academic and research computing and computational science communities, and seeks to illustrate the viability of these services as an option for leading-edge research across a broad scope of science. The project helps researchers understand the potential benefit of larger-scale commercial platforms for scientific application workflows such as those currently using NSF’s High-Performance Computing (HPC). It also explores how scientific workflows can innovatively leverage advancements provided by commercial cloud providers. The project aims to accelerate scientific discovery through integration and optimization of commercial cloud service advancements; identify gaps between cloud provider capabilities and their potential for enhancing academic research; and provide initial steps in documenting emerging tools and leading deployment practices to share with the community.

Additionally, a recent example of a partnership between Internet2 and government is the NSF Cyberinfrastructure Center of Excellence Demonstration Pilot (Award #2137123) awarded to

Internet2 to support the Minority Serving Cyberinfrastructure Consortium (“MS-CC”). This Center of Excellence (“CoE”) Demonstration Pilot will initially support the advancement of research cyberinfrastructure capabilities and support systems for Historically Black Colleges and Universities (“HBCUs”) and Tribal Colleges and Universities (“TCUs”), with the goal of expanding to Hispanic Serving Institutions (“HSIs”) and other Minority Serving Institutions (“MSIs”). Internet2 will serve in a facilitating role by coordinating efforts that will allow participating organizations to work together to address common challenges that may be more difficult for each to resolve independently. The abstract states the following:

A key outcome of this grant is the formalization of a vibrant community of practice across MS-CC campuses that involves collaboration on cyberinfrastructure, education, and research applications. This grant enables HBCUs, TCUs, HSIs and other MSIs to accomplish together what they cannot do separately. The MS-CC is broadening participation in science, technology, engineering, and mathematics (STEM) by historically underrepresented groups in the United States’ research enterprise, enabling new perspectives to emerge and expand capabilities for the nation. MS-CC is advancing our nation’s economic growth, national security and global prosperity in ways that reflect the unique expertise and talent from HBCUs, TCUs, HSIs and other MSIs.

There also are examples of public-private partnerships between Internet2, regional R&E networks, and the National Institutes of Health (“NIH”) related to high-speed networking for the dissemination of NIH data, secure access controls for NIH’s Electronic Research Administration Portal, and NIH’s Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability ([STRIDES](#)) Initiative, which “allows NIH to explore the use of cloud environments to streamline NIH data use by partnering with commercial providers.”

Internet2, together with Mid-Atlantic Crossroads (MAX), the regional R&E network led by the University of Maryland, provides research network capacity to NIHnet (NIH’s backbone network). As described by NIH (<https://www.cit.nih.gov/services/networks>):

NIHnet is a high-speed, highly available network that interconnects NIH, the commodity internet, and the *Internet2 research network*. NIHnet runs a 100G backbone – which runs up to 50,000 times faster than the Internet access for most U.S. households – that provides fast, secure, consistent connections for network traffic, even when there are hardware or software disruptions or slowdowns.

Internet2 also is working with NIH to implement its upcoming requirements for multi-factor authentication (“MFA”) and identity proofing levels in connection with federated access to many of NIH’s services. NIH’s Electronic Research Administration (“eRA”) Portal began requiring

users to sign in with MFA on September 15, 2021. This change affects everyone who accesses eRA. Internet2 prepared an eRA Readiness Guide to IdP operators in the InCommon Federation to understand how this change impacts campus users. It also helps campuses prepare support options to assist users during the transition (<https://spaces.at.internet2.edu/display/federation/get-nih-ready>).

In support of NIH’s Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (“[STRIDES](#)”) Initiative, Internet2 partnered with NIH and Google Cloud Platform (“GCP”) to allow NIH-funded researchers to leverage Internet2’s NET+ GCP terms to use a single agreement to access GCP for their enterprise and research needs as part of the initiative. Internet’s NET+ program provides custom terms of service for cloud services developed by Internet2 higher education members specifically to meet the unique needs of research and higher education institutions.

All of these cited examples demonstrate a long track record of successful public-private partnerships between regional and national R&E networks, government, academia, and the for-profit private sector that can be used as models for many of the NAIRR capabilities outlined in Topic D.

6. Where do you see limitations in the ability of the NAIRR to democratize access to AI R&D? And how could these limitations be overcome?

The first step should be to define “democratize.” In the broadest sense, democratization should include providing the NAIRR access to all higher education institutions and associated organizations engaging ML/AI research or applying ML/AI across scientific disciplines. In this context, the limitations to democratize access must include providing the fundamental enabling cyberinfrastructure to those institutions. One envisioning of overcoming these limitations can be found in a recent R&E community-developed paper, “The Minds We Need” (<https://mindsweneed.org/>), which outlines three core actions:

- **Connect every community college, every MSI, and every college and university, including all urban, rural, and tribal institutions** to a world-class and secure R&E infrastructure, with particular attention to institutions that have been chronically underserved;
- **Engage and empower every student and researcher** everywhere with the opportunity to join collaborative environments of the future, recognizing that the “last mind connected may be the mind we need;” and

- **Ensure American competitiveness and leadership** by investing holistically in national R&E infrastructure as a sustainable system.

As highlighted above, particular attention should be paid to institutions that have been chronically underserved. Those institutions have additional challenges, some of which can be found in a recent stakeholder alignment survey conducted by Internet2 and the MS-CC to better understand the needs of HBCUs, TCUs, and HSIs (<https://internet2.edu/solutions/minority-serving-institutions/>).

A key objective of this effort is to help MSIs identify the science, engineering, health, social science, and humanities education and research priorities that call for increased access to and use of data management and computing resources by the higher education community. It is clear that many of the key findings of the survey would apply to the democratization of access to AI R&D.

Key findings of needs identified include:

- **Basic needs:** Across HBCUs, HSIs, and TCUs there is a deep need for basic infrastructure support, such as broadband Wi-Fi on campus and at home for students, staff, and faculty (heightened by the pandemic).
- **Consistency across institutions:** Although there are some unique considerations for certain types of institutions (such as data sovereignty with TCUs), the vast majority of HBCUs, HSIs, and TCUs have similar responses.
- **Workforce development:** Students need literacy and advanced skills with data and computing (the educational mission), while faculty and staff need training and support for a robust cyberinfrastructure (the research mission).
- **Collaboration:** There is strong support for collaboration across institutions to accomplish together what they cannot do separately (with little support for each acting on their own).
- **Institutional operations:** Administrators need a more accessible and responsive data infrastructure for campus operations, surfaced when asked about research data and computing.
- **Societal impact:** There is strong potential for data and computing to advance research on issues central to community culture and disparities in society, with infrastructure as a constraint on achieving these broader impacts.

The first step in overcoming these limitations is to include representatives from underserved institutions and communities from the start of the process.

CONCLUSION

For the foregoing reasons, Internet2 respectfully requests that NSF and OSTP will consider these recommendations and looks forward to working together on this important initiative.

Respectfully submitted,

/s/ Howard Pfeffer

Howard Pfeffer
President and CEO