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

Deloitte.

Implementing a National Artificial Intelligence Research Resource

October 1, 2021

In response to the Implementation Plan for a National Artificial Intelligence Research Resource Request for Information (RFI)

RE: RFI Response: National AI Research Resource

Deloitte Consulting LLP (Deloitte) is pleased to submit to the Office of Science and Technology Policy (OSTP) and National Science Foundation's (NSF) Request for Information to implement the National Artificial Intelligence Research Resource (NAIRR) Implementation Roadmap.

Deloitte believes the NAIRR initiative fills a critical and strategic gap in the government's AI initiative. To support its success and address all RFI parameters, our response draws on Deloitte's AI expertise in both the government and private sector domain. This includes our in-house data science talent, experts across AI and government, experience developing AI solutions for public and private sector clients, and familiarity with the perspectives of AI researchers through regular surveys and our professional services.

We remain ready and interested to support the NAIRR Task Force (TF). We would be pleased to share our experience advising and implementing AI across government and commercial with OSTP and NSF. Should you have any questions, please contact me at ([redacted].

Sincerely,

Ed Van Buren
Principal
AI in Government Leader
Deloitte Consulting LLP

Table of Contents

Table of Contents

Company Profile

Company Name	Deloitte Consulting LLP
Headquarters Location	New York City
Contact Name	Ed Van Buren
Contact Title	Principal, Deloitte Consulting LLP
Contact Email Address	[redacted]
Contact Phone Number	[redacted]
Primary Type of Service(s) Provided	Software Development, Professional Services, Management Consulting, Technical Support, Maintenance, and Support Services

Executive Summary

Artificial intelligence holds great promise for U.S. economic growth and prosperity as well as national competitiveness, more generally. However, the evolution of AI has been inhibited by issues of data access, security, and quality; the increasing computational demands and complexity of AI modeling; and differing interests and disparate resources of various AI constituencies in the private sector, government, and academia.

These issues have routinely emerged across industries. In our annual State of AI in the Enterprise survey, we observed varying rates of AI adoption, varying preferences in AI development approaches, and varying competencies for effective AI procurement between private sector and public sector respondents. We recently developed an AI Dossier to help leaders understand when, where, and how to deploy AI within their organization. The Deloitte AI Institute for Government collaborates with an ecosystem of academic institutions and thought leaders to identify practical ways to develop and deploy AI.

While our response details considerations and pathways for the establishment and sustainment of the NAIRR, we highlight three areas the TF will need to navigate: 1) need for openness & transparency to foster cutting-edge innovation while adhering to security demands of sensitive

and proprietary data; 2) need to collaboratively address challenges of interoperability across data types and models; and 3) need to foster rapid innovation and R&D with a focus on ethical and trustworthy AI as well as diversification of AI developer talent.

As detailed below, we believe the solution lies within the NAIRR governance model, its resource base, and stakeholder inclusion so AI research is not misunderstood and is widely beneficial.

Responses to RFI Questions

Q1: What options should the Task Force consider for any of roadmap elements A through I above, and why? [Please take care to annotate your responses to this question by indicating the letter(s) of the item (A through I in the list above) for which you are identifying options.]

	Response
A.	The TF should prioritize a flexible development process, and **employ an agile methodology**, allowing lessons learned to be incorporated iteratively. Defining success criteria from the onset will be critical to a successful implementation. The TF should **consider Key Performance Indicators (KPIs) based on goals** of the NAIRR and develop yearly metrics targets, even if preliminary and directional, to drive and evaluate success. Illustrative metrics could include: number of projects executed, number and variety of quality dataset uploads and downloads, activity levels on knowledge portals, computational resource availability and sharing, new users onboarded, and partnerships established. Additionally, the NAIRR can capture demographic data to assess if underserved communities, institutions, and regions are accessing and utilizing the NAIRR.
B. & C.	The NAIRR's success will depend on the organization responsible for program implementation, deployment, and management. Considerations for the TF during the selection process could include: • **Mission Alignment**: The goal of "democratizing access to the cyberinfrastructure that fuels AI research and development" should fit the organization's mission - this facilitates shared purpose and impact. • **Positioning**: The organization should have industry agnostic scope to allow research across industry, government, and academia. An agnostic approach will help with and contribute to an equitable partner experience. • **Authority**: The organization should have resources and capacity to establish direction, make programmatic decisions, and manage resources. To maximize adoption and utilize existing resources and expertise, the NAIRR should consider a **hub and spoke** governance model. This would centralize the NAIRR as a **hub** that governs the standup and drives policy, standards, and driving adoption, while coordinating with existing agency missions, partnerships, engagements, infrastructures, and research communities (**spokes**). As the central hub, the NAIRR could establish a PMO charged with strategic direction, programmatic decision making, and resource allocation. This approach allows each research spoke to be established based on individual requirements and provides them autonomy which can lead to innovation. The NSF AI

	Institutes Program is a good example of hub and spoke governance where the NSF is a central hub and the partnered universities are the spokes.
D.	To democratize AI R&D capabilities, the TF will need to prioritize foundational capabilities available in the initial operating capability. This includes: • Data infrastructure with pretrained models and quality datasets • Hybrid-cloud compute infrastructure, compute power and services • Knowledge management portal for educational tools and services • Proper testing & evaluation capabilities to allow models to be validated • Change Management and strategic comms to drive engagement and adoption Please see 3.2 for a full response.
E.	The ability to share and disseminate high quality data across the NAIRR is contingent on the successful mitigation of several operational and technical concerns. Each will require discrete risk mitigation techniques to overcome the concerns. • **Security Concerns** – The NAIRR should create a framework to detect security and cyber risks and develop incident response plans in the event of a compromise. • **Usage Concerns** – To avoid data confusion / misinterpretation, each dataset should come with an abstract or documentation to describe the data and its formatting. • **Attribution Concerns** – License requirements for data uploaders to populate can be combined with an auto-attribution feature from the uploader's profile to ensure that if institutions reuse datasets / resources, they cannot remove the license or text that is copyrighted. • **Cultural Barriers** – Siloed mindsets can inhibit progress in research, data, and models. To encourage data sharing, institutions can be rewarded and credited with additional resources (i.e., storage, compute, etc.) for their contributions.
F.	To successfully operate and house research from various research hubs, the program must **invest in state-of-the-art technologies** to encrypt data and manage access to its resources, to include data, models, and related research. To identify appropriate security and access management controls, the NAIRR could work with National Institute of Standards and Technology (NIST) to curate standards to address security and operational constraints (i.e., NIST 800-53). This will help the NAIRR develop a resilient **federated Identity and Access Management (IdAM) control system** across its hybrid-cloud infrastructure. To enable this, the NAIRR should consider a **zero-trust architecture**, a data-centric security framework which requires strict identity verification for devices and personnel. An approach of least privilege can enhance the NAIRR's cyber risk posture by limiting impact of cyber breaches and improving containment capabilities. This approach will require the orchestration of multiple organizational components, including buy-in from oversight organizations, capability providers, and research hubs.
G.	**Trust** between users and developers is key in the discipline of AI. Trust is earned over time and lost in an instant. In response to Executive Order 13859, NIST developed a framework to establish standards to bolster public trust and confidence in AI applications. These standards address societal and ethical issues, governance, and privacy.

	To assess privacy and civil rights requirements associated with the NAIRR conducted research, the NAIRR should consider a framework to assist researchers with assessing elements underpinning this trust, including data attributable to consumers, citizens, or patients. The NAIRR should consider **NIST's Trustworthy AI Framework**.
H.	Sustained and predictable **federal funding** will be essential for the NAIRR's initiation. This will allow objectivity and independence necessary to seed / support cutting-edge innovations and access to educational resources and data. Alternative arrangements, such as subscription-based or fee-for-services may undermine the NAIRR's ability to develop an active user base, promote innovation, or skew / narrow direction of R&D priorities. **Private sector partnerships** remain a potentially impactful multiplier for the NAIRR and must navigate unique interests. Prospective private sector partners will be incentivized to engage in NAIRR initiatives if they can showcase / develop internal talent, permit cooperative problem solving with public sector researchers, highlight their tools and services, or provide compute resources and technology to funded R&D activities, potentially at discounted rates if it provides them visibility and branding opportunities.
I.	The NAIRR should be established in phases, with iterations to assess, evaluate, and adjust. Stakeholders should agree on capabilities needed for success. Parameters for establishment and sustainment can be broken into three steps 1) **determine public authority** (legal framework and political realities) 2) **define project needs & objectives** (i.e., speed, efficiency, degree of certainty) and 3) **determine owner for each project** (i.e., capabilities, financial, operational, and risk transfer). To help the NAIRR sustain its operations, the TF can consider an executive steering group, similar to the one operated by the Joint AI Center (JAIC). While the JAIC's original roles and responsibilities were derived from the National Defense Authorization Act, the Agency established an executive steering group to enable stakeholders to interact with one another and align priorities. This group is composed of senior leaders and General officers, who work across areas (i.e., acquisition, workforce development, standards, AI ethics, and AI policy) to ensure conversations and perspectives are integrated across the enterprise.

Q2: Which capabilities and services (see, for example, item D above) provided through the NAIRR should be prioritized?

One of the key capabilities and services the NAIRR can provide to users is **high quality clean datasets and pretrained models**. Providing resources such as clean datasets and pre-trained models allow users to quickly spin up models which can be iterated on. Availability and access to data is one of the greatest obstacles. The NAIRR should prioritize user-friendly methods to help researchers and institutions identify datasets best suited for their problems. This involves data categorization (supervised, unsupervised, cleaned, uncleaned, etc.) to start, and can be advanced into a recommender system to improve information sharing and reduce friction for finding available datasets. One example is NIH's FAIRshake program which makes data findable, accessible, interoperable, and reusable. Furthermore, a repository of trained models shared between institutions allows for a model economy that can scale by building on each other.

With the development of deep learning models and others requiring increasing compute, the need for a **shared compute infrastructure** should be prioritized. Shared infrastructure with robust compute power is critical to the progression of AI. The easiest solution to provide compute power is established cloud and hardware vendors. To begin, the NAIRR should instantiate a single CSP, with the goal of building a hybrid on-premise and multi-cloud environment, which allows for greater interoperability. This approach reduces initial costs and programmatic complexity and allows for at-speed scalability. There are instances which could require localized processing and high-performance computers. This on-premise capability could be used for certain use cases, datasets and / or model development that would benefit from high performance compute.

The NAIRR should invest in a **knowledge management portal** to provide educational tools for new and seasoned data scientists. This portal can foster collaboration among institutes and researchers and create an open knowledge network to aid in idea creation and problem solving.

To ensure research conducted through the NAIRR's resources does not introduce bias and risk, **testing capabilities and procedures** must be in place for users when evaluating and validating their models. These tools should be flexible and allow for continuous testing, integration, and deployment, which will instill justified confidence in models produced and used.

For the NAIRR to sustain its operations, the TF could implement a robust **change management plan** and ensure **strategic communications** drive engagement and adoption amongst stakeholders. Throughout NAIRR's operation, it is critical provisioning and monitoring of devices are done on a continual and automated basis to prevent data leaks and unauthorized use.

Q3: How can the NAIRR and its components reinforce principles of ethical and responsible research and development of AI, such as those concerning issues of racial and gender equity, fairness, bias, civil rights, transparency, and accountability?

Ethical and Trustworthiness evaluation criteria need to be part of the process when using the NAIRR and should be instantiated at the beginning of a research project. Anytime there is model development and training, there is a risk of biases being implicitly introduced. The NAIRR should consider how to mitigate risks from several types of biases, to include **data-driven bias** which can result from the data used to train the model, **confirmation bias** where models filter out relevant information from the user, as they think it is "noise", or **judgement bias**, where models can make determinations based on ambiguous data.

Any model developed needs to be evaluated to determine it meets the elements of Trustworthy AI. Rather than manual reviews, the NAIRR can develop an assurance framework to support the integrity of models and underlying data. Additionally, the NAIRR **could establish a quality review process which assesses certain properties of the data, models, and questions being answered.** For example, the review process could include acknowledging legal and

privacy restrictions, data consistency across population cohorts, data governance rules, and an assurance that datasets are representative of the issue the model is trying to solve.

Example: The NAIRR can draw from financial institutions that put algorithm risk into practice at scale. After the global financial crisis exposed the risks of inaccurate algorithm-driven models, the Federal Reserve and Office of the Comptroller of the Currency (OCC) issued the Supervisory Guidance on Model Risk Management (SR 11-7). It required identification and estimation of adverse consequences of inaccurate or misused models. SR 11-7 obligates users understand the limitations of models and avoid using models for uses other than originally intended. It mentions models should be validated regularly to ensure they are performing as expected.

Q4: What building blocks already exist for the NAIRR, in terms of government, academic, or private-sector activities, resources, and services?

Government

- **Data.gov:** GSA's Technology Transformation Services manages a repository of metadata, data.gov. The site stores 320+ datasets and resources such as tools, policies, and case-studies. Data.gov and other repositories such as Data USA could be leveraged to enhance data access for the NAIRR.
- **NIH Data Commons:** NIH ran a pilot for NIH Data Commons from 2017-2018. The pilot tested the best ways to build and implement a cloud-based platform designed to store, share, access and interact with data and other digital objects. For example, a tool called **FAIRshake** evaluates how easy / difficult it is to find, access, interoperate, and reuse digital objects. There are similar types of programs that exist, such as Dept. of Transportation's Safety Data Commons and USAID's Development Data Commons.
- **JAIC JCF:** JAIC is building an AI data and algorithm development platform called JCF. The NAIRR could establish a data exchange with JCF to bolster their data catalogue.

Academic

- **Colleges and Universities:** Schools across the country are standing up AI Institutes, and programs the NAIRR can tap into through partnerships. The University of California San Diego, along with five other universities, is trying to address the challenge of scaling across several areas including health, semiconductor chip design, robotics, and networks. In addition, numerous other university consortiums are looking to conduct research, such as the University of Chicago with the Digital Transformation Institute and Howard University, which is part of an Historically Black Colleges and Universities (HBCU) consortium developing curriculums and talent in quantum computing.

Private Sector

- **Computational Resources:** CSPs offer elastic compute and data storage capabilities ideal for a holistic computing ecosystem. By leveraging a hybrid cloud, the NAIRR can reap the benefits of cloud technology while still leveraging existing on-premises compute

infrastructure owned by research hubs. In a hybrid model, NAIRR research hubs will have access to pre-built AI tools that support their research.

- **Private Sector Investments / Products:** The NAIRR can explore existing private sector assets to jump start the creation of the NAIRR with pre-configured capabilities. An example is Deloitte's Cortex AI, which draws from Deloitte use cases, solutions, and data, and acts as an accelerator to AI adoption by applying them to help clients devise intuitive solutions, scale adoption faster, and develop a competitive edge in their work.

Q5: What role should public-private partnerships play in the NAIRR? What exemplars could be used as a model?

Public Private Partnerships (PPP) can be a powerful tool in tackling complex public service programs such as the NAIRR. PPPs can help government engage diverse stakeholders, control costs, and mitigate risks. However, **PPPs are not a one-size-fits-all**. In order to choose the right PPP model, government must determine the project components, public authority, goals, and the best "owner" for each phase from among the various service providers, vendors, system integrators, and startups that comprise the stakeholder ecosystem.

There are numerous models of PPP tailored to uses ranging from building highways to managing public goods such as airports or transportation infrastructure. PPPs have even been used to shape digital infrastructure in a similar way to what the NAIRR is trying to accomplish.

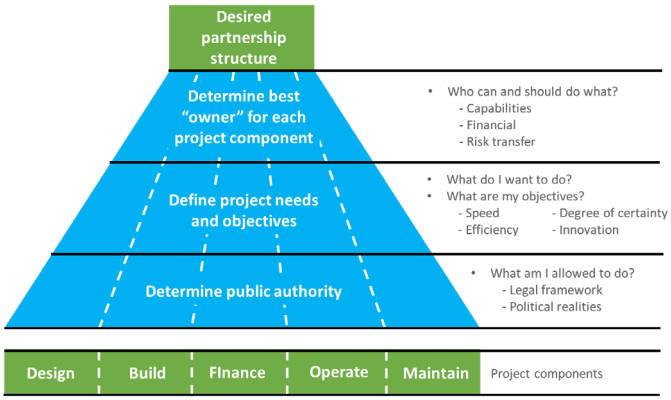

Source: Deloitte analysis

For example, NIH used PPP models to create NIH Data Commons, a cloud-based platform where researchers could store, share, access, and use digital files generated from biomedical research. The NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability Initiative widened the lens of partners, working with CSPs to allow 2,500 NIH-funded institutions to explore cloud and **ML capabilities to generate, analyze, and share data**. Recently, data has shown that these PPPs can work quickly and on sensitive data, creating a centralized, secure, and **cloud-enabled data platform** to analyze real-world COVID-19 patient data across government, hospital systems, and research institutions.

Q6: 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?

Data quality, management, and accessibility limitations
There are operational and technical barriers for the NAIRR to fulfill its goal of democratizing access to AI R&D. Operationally, success of the NAIRR will be heavily contingent on the platform's management of and access to quality datasets for users. As is often said in data science, "*garbage in garbage out"*, meaning bad data results in bad outputs. The NAIRR will need to gain / maintain the trust of users and develop mechanisms to encourage data sharing. For example, users from the scientific community may use journal standards to dictate whether submission of replication datasets / code are required. A solution is to develop data quality standards each research hub is encouraged to meet as a part of the NAIRR agreement. This could be incentivized through access to storage and / or additional cloud compute resources.

To tackle technical challenges, data cataloguing across the platform can help users quickly index large volumes of data to find relevant datasets. The NAIRR can offer templates and scripts to fix issues like formatting, missing values, and irregularities to help foster quality data. The NAIRR could survey data housing, assembly, cataloguing, and sharing practices across disciplines as a step toward establishing data quality standards that meet the needs of all user groups.

Cybersecurity Threats
There has been exponential growth in cyberattacks the past couple years, and for a program reliant on a decentralized governance model, it is imperative security mechanisms are implemented from bottom up to maintain program integrity. Platforms that house data and provide computational resources are charged with prioritizing confidentiality, integrity, and availability of data. For this, the NAIRR could leverage software to intelligently manage IdAM controls across the enterprise. The NAIRR could invest in tools to continuously monitor the network to increase visibility into network activity to detect potential security breaches, allowing the incident response team enough time to mitigate risks.

Organizational Culture
Organizations may be fearful or resistant to AI across their enterprise. Job displacement and ethical implications may dissuade users. The NAIRR can prepare R&D initiatives by emphasizing stakeholder inclusion, sensitivity to interests of data subjects and modeled outcomes. The NAIRR should incorporate a change management plan inclusive of all stakeholders and deliver strategic comms regularly to avoid misunderstandings or overlooked concerns and biases.

Diversity in AI
The blossom of AI has not yet been distributed in an equitable manner. Underrepresented groups have been locked out from both a talent (researchers and developers) and as consumers / beneficiaries of AI R&D. To help combat this, the NAIRR could consider mechanisms to prioritize allocated resources or research priorities that will benefit small startups and / or underrepresented institutions such as HBCUs.