

Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem

*An Implementation Plan for a
National Artificial Intelligence Research Resource*



January 2023

**Strengthening and Democratizing the U.S.
Artificial Intelligence Innovation Ecosystem:**
*An Implementation Plan for a National
Artificial Intelligence Research Resource*

National Artificial Intelligence Research Resource Task Force

January 2023

Dear Mr. President and Members of Congress,

Artificial Intelligence (AI) is changing our country and our world. From how citizens navigate their daily lives to how researchers drive discoveries in the lab to how manufacturers build products, AI is giving rise to new capabilities. New AI and AI-driven discoveries and capabilities hold the potential to drive practical solutions to address critical global challenges such as food production, climate change, poverty, and cancer. We have only started to scratch the surface of what is possible, and cannot afford to miss out on seizing the opportunity for leveraging AI to serve the public good.

However, the opportunities to pursue cutting-edge AI research and apply AI to new domains and challenges are currently not accessible by all of America's incredible talent nor harnessed by the public sector. Much of today's AI research relies on access to large volumes of data and advanced computational power, which are often unavailable to researchers beyond those at well-resourced technology companies and universities. This access divide limits the ability to leverage AI to tackle the big challenges in our society. It also constrains the diversity of researchers in the field and the breadth of ideas incorporated into AI innovations, contributing to embedded biases and other systemic inequalities found in AI systems today.

Recognizing this challenge, in the National AI Initiative Act of 2020, Congress directed the National Science Foundation (NSF), in consultation with the White House Office of Science and Technology Policy (OSTP), to establish a task force to create a roadmap for a National AI Research Resource (NAIRR)—a shared research infrastructure that would provide AI researchers and students with significantly expanded access to computational resources, high-quality data, educational tools, and user support.

This final report of the NAIRR Task Force presents a roadmap and implementation plan for a national cyberinfrastructure aimed at overcoming the access divide, reaping the benefits of greater brainpower and more diverse perspectives and experiences applied to developing the future of AI technology and its role in our society. Such a national cyberinfrastructure also presents a unique and critical opportunity to "design in" the standards for responsible AI research practices and governance processes that uphold our priority to develop and harness these groundbreaking technologies in a manner that reinforces our Nation's democratic values and Americans' personal freedoms.

OSTP and NSF formally launched the NAIRR Task Force in June 2021, appointing 12 leading experts equally representing academia, government, and private organizations. Over the course of its work, the Task Force held 11 public meetings, engaged with 65 experts on a wide range of aspects related to the design of the NAIRR, and considered responses from the public to two requests for information. We extend our gratitude to the members of the Task Force who have donated an extraordinary number of hours of their time to this effort, as well as to the many members of the public who have contributed their expertise and provided inputs to the Task Force. The result of the last one and one-half years of effort is this final report.

We see the NAIRR as a foundational investment that would amplify efforts across the Federal Government to cultivate AI innovation and advance trustworthy AI. Research, experimentation, and innovation are integral to our progress as a Nation, and it is imperative that we engage people from every zip code and every background to live up to America's unique promise of possibility and ensure our leadership on the world stage.

The work of the NAIRR Task Force and this report will be an invaluable resource as we work collaboratively across government and across sectors to drive this important work forward.

Sincerely,



Sethuraman Panchanathan
Director
National Science Foundation



Arati Prabhakar
Assistant to the President for
Science and Technology
Director, Office of Science and
Technology Policy

Executive Summary

Artificial Intelligence (AI) is an engine of innovation that is driving scientific discovery and economic growth. It is increasingly becoming an integral component of solutions that stand to impact everything from routine daily tasks to societal-level challenges, ultimately serving the public good. At the same time, there are also concerns that AI could have negative social and environmental consequences. To realize the positive and transformative potential of AI, it is imperative to harness all of America's ingenuity to advance the field in a manner that addresses societal challenges, works for all Americans, and upholds our democratic values.

Yet progress at the current frontiers of AI is often tied to access to large amounts of computational power and data. Such access today is too often limited to those in well-resourced organizations. This large and growing resource divide has the potential to limit and adversely skew our AI research ecosystem. The imbalance threatens our Nation's ability to cultivate an AI research community and workforce that reflect America's rich diversity and the ability to harness AI to advance the public good.

A widely accessible AI research cyberinfrastructure that brings together computational resources, data, testbeds, algorithms, software, services, networks, and expertise, as described in this report, would help to democratize the AI research and development (R&D) landscape in the United States for the benefit of all. It would help create pathways to broaden the range of researchers involved in AI, and to grow and diversify approaches to, and applications of, AI. This cyberinfrastructure can also help to open up new opportunities for progress across all scientific fields and disciplines, including in critical areas such as AI auditing, testing and evaluation, trustworthy AI, bias mitigation, and AI safety. Increased access and a diversity of perspectives can, in turn, lead to new ideas that would not otherwise materialize and set the conditions for developing AI systems that are inclusive by design.

As part of the National Artificial Intelligence Initiative Act of 2020, Congress established the National Artificial Intelligence Research Resource (NAIRR) Task Force to "investigate the feasibility and advisability of developing" the NAIRR as a national AI research cyberinfrastructure, and "to propose a roadmap detailing [how the NAIRR] should be established and sustained." The recent CHIPS and Science Act of 2022 reinforces the importance of democratizing access to a national AI research cyberinfrastructure, via investments that will accelerate development of advanced computing—from next-generation graphics processing units to high-density memory chips—as well as steps to actively engage broad and diverse U.S. talent in frontier science and engineering, including AI.

This final report is the culmination of the Task Force's 18-month effort to develop a vision and implementation plan for establishing the NAIRR. It builds on the findings and recommendations outlined in the Task Force's interim report released in May 2022, providing an implementation plan to achieve the objective of the NAIRR: to strengthen and democratize the U.S. AI innovation ecosystem in a way that protects privacy, civil rights, and civil liberties.

The NAIRR should be established with four measurable goals in mind, namely to (1) spur innovation, (2) increase diversity of talent, (3) improve capacity, and (4) advance trustworthy AI. The NAIRR should meet these goals by supporting the needs of researchers and students from diverse backgrounds who are pursuing foundational, use-inspired, and translational AI research. These users should be U.S.-based or affiliated with U.S. organizations, to include academic institutions, non-profit organizations, and startups or small businesses.

The NAIRR should comprise a federated set of computational, data, testbed, and software resources from a variety of providers, along with technical support and training, to meet the needs of this target user base. The specific design, implementation, and evaluation of the NAIRR should be centered around the four key goals, and should support the collection of data for assessment of key indicators of system performance and success in progress toward these goals.

The NAIRR administration and governance should follow a cooperative stewardship model, whereby a single Federal agency serves as the administrative home for NAIRR operations and a Steering Committee comprising principals from Federal agencies with equities in AI research drives the strategic direction of the NAIRR. A Program Management Office within the administrative home agency should provide funding and oversight for an independent Operating Entity that manages the day-to-day operations of the NAIRR. The Steering Committee, co-chaired by the National AI Initiative Office (NAIIO), would incorporate interests and perspectives from across Federal agencies in the governance of the NAIRR. These agencies should also directly support resource providers whose resources, in federation, would constitute the NAIRR. Diverse perspectives and expertise should be tapped to inform the NAIRR's operations through a User Committee, a Science Advisory Board, a Technology Advisory Board, and an Ethics Advisory Board that provide advice to the Operating Entity.

The NAIRR should provide access to a federated mix of computational and data resources, testbeds, software and testing tools, and user support services via an integrated portal. Computational resources should include conventional servers, computing clusters, high-performance computing, and cloud computing, and should support access to edge computing resources and testbeds for AI R&D. Open and protected data should be made available under tiered-access protocols and co-located with computational resources. The Operating Entity should not itself operate the totality of the computer hardware that composes the NAIRR; instead, computing, along with data, testing, and training resources, should be delivered as services by partner resource providers selected through Federal agency or multi-agency funding opportunities. When fully implemented, the NAIRR should address both the capacity (ability to support a large number of users) and capability (ability to train resource-intensive AI models) needs of the AI research community.

The NAIRR must be broadly accessible to a range of users and provide a platform that can be used for educational and community-building activities in order to lower the barriers to participation in the AI research ecosystem and increase the diversity of AI researchers. The NAIRR access portal and public website should provide catalogs and search and discovery tools to facilitate access to data, testbeds, and educational and training resources serving a range of experience levels.

The NAIRR should set the standard for responsible AI research through the design and implementation of its governance processes. The NAIRR must be proactive in addressing privacy, civil rights, and civil liberties issues by integrating appropriate technical controls, policies, and governance mechanisms from its outset. The Operating Entity should work with its Ethics Advisory Board to develop criteria and mechanisms for evaluating proposed research and resources for inclusion in the NAIRR from a privacy, civil rights, and civil liberties perspective. Regular training should be required to build NAIRR users' awareness about rights, responsibilities, and best practices related to privacy, civil rights, and civil liberties in AI research, in accordance with the Blueprint for an AI Bill of Rights published by the White House Office of Science and Technology Policy in October 2022.

The NAIRR should implement system safeguards in accordance with established guidelines. These guidelines include those developed by the National Institute of Standards and Technology (NIST) and the Five Safes framework: safe projects, safe people, safe settings, safe data, and safe outputs. The Operating Entity should design the NAIRR cyberinfrastructure to consist of multiple tiers, starting with two primary zones: an open science zone "NAIRR-Open" and a secure zone "NAIRR-Secure." Each zone should federate computational, network, and data resources operating in accordance with security and access-control policies that are uniform within the zone, but different between zones, reflecting the different priorities and needs of the users and resource operators. NAIRR-Open should adopt the best practices developed over two decades in the open science community; be consistent with Federal open data, open government, and research security policies; and manage access using single sign-on authentication and a resource allocation mechanism managed by the Operating Entity. NAIRR-Secure should consist of one or more secure enclaves adhering to a common set of security controls, and have the ability to support security requirements arising from legally protected data.

NAIRR implementation should occur over four phases, beginning immediately after the publication of this report. In phase one, Congress should authorize and appropriate funds to establish the NAIRR. The administrative home agency and the NAIIO should coordinate the formation of the Steering Committee and stand up a Program Management Office, which will then prepare the solicitation for the Operating Entity and manage the selection process.



Phased NAIRR Implementation Timeline

In phase two, the Operating Entity should establish its activities and oversee creation of the NAIRR portal and user interface, building in appropriate technical and policy controls. The architecture should support collection of key performance indicators for evaluation of NAIRR progress. Resource providers should be selected via coordinated, multi-agency funding opportunities ideally released within six months of the initial Operating Entity award.

In phase three, the NAIRR should achieve initial operational capability and the Operating Entity should also formalize the policies, processes, and initial technical resources to be made available to AI researchers. Initial capabilities include (1) a portal and user support resources, (2) a mix of computational resource providers, (3) an allocation and identity system, and (4) a data publication system. In phase four, activities should transition from building out the NAIRR to establishing steady-state operations, as well as the planned evolution of NAIRR resources in response to user uptake and demand.

Finally, the Task Force also presents a pilot option for the implementation of the NAIRR that would be initiated in parallel with the above phases to expedite the availability of NAIRR resources to the AI R&D community.

As envisioned, the impact of the NAIRR will be significant and far-reaching, enabling researchers to tackle problems that range from routine tasks to global challenges. In order to achieve its vision and goals, the Task Force estimates the budget for the NAIRR as \$2.6 billion over an initial six-year period. The bulk of this investment (\$2.25 billion) is to fund the resources to be made accessible via the NAIRR, through appropriations to multiple Federal agencies. The Task Force estimated this budget based on recent costs of advanced computing resources as well as data, training, and software resources; estimates of usage levels to meet the current needs of the AI R&D community; and expected growth of the AI R&D community. Resource providers should be brought online every two years with a six-year lifetime, so that a new \$750 million investment is made every two years to ensure that the NAIRR resources remain state-of-the-art. The Operating Entity will require between \$55 million and \$65 million per year to support the coordination and management of NAIRR activities. An additional \$5 million per year is budgeted for external evaluation of the Operating Entity and NAIRR performance.

The vision for the NAIRR laid out in this report is designed to meet the national need for increased access to the state-of-the-art resources that fuel AI innovation. The roadmap for achieving this vision builds on existing Federal investments; designs in protections for privacy, civil rights, and civil liberties; and promotes diversity and equitable access. **If successful, the National AI Research Resource would transform the U.S. national AI research ecosystem and facilitate the ability to address societal-level problems by strengthening and democratizing participation in foundational, use-inspired, and translational AI R&D in the United States.**

Contents

1. Introduction	1
The Current Landscape of AI R&D	1
An Opportunity for Strengthening AI R&D in the United States	3
The National AI Research Resource Task Force	4
Structure of This Report	6
2. A National Cyberinfrastructure to Democratize and Accelerate AI R&D	7
NAIRR Vision and Goals	7
The NAIRR User Base	8
NAIRR Constituents	10
<i>Government</i>	10
<i>Academia</i>	10
<i>Industry</i>	11
<i>Civil Society</i>	11
3. NAIRR Organization, Management, and Governance	12
NAIRR Organizational Structure	12
<i>Steering Committee</i>	14
<i>Individual Agencies</i>	15
<i>Program Management Office</i>	17
<i>Operating Entity</i>	18
<i>NAIRR Staff and Executive Leadership Team</i>	20
<i>NAIRR Advisory Boards</i>	21
<i>Evaluation Entity</i>	22
User Access and Resource Allocation	22
Privacy, Civil Rights, and Civil Liberties Protections	24
Scientific Integrity	26
System Security and User Access Controls	26
<i>Open-Source Principles</i>	28
Environmental Sustainability	28
4. NAIRR Structure and Specifications for Resource Elements	30
Access Portal and User Interface	30
Computational Resources	31
<i>Capacity and Capability</i>	31
<i>NAIRR Software Resources</i>	31
Data and Datasets	32
<i>Dataset Acceptance Criteria and Metadata Standards</i>	33
<i>Role of the Operating Entity in Incentivizing and Curating Contributed Datasets and Other Resources</i>	33
<i>The NAIRR and Existing Federal Government Data</i>	34
<i>Legal Compliance</i>	35
Co-Location of Resources	36
Educational Tools and Services	36

Tiered Technical Training and Support	37
Curation of Training Materials	37
Platform for Educational Activities	37
Technical Integration	37
Software for Integration	37
Integrating Data Resources	38
Testbeds	38
5. Phased Buildout of NAIRR Organization and Resources	40
Phase 1: Program Launch and Operating Entity Selection	41
Phase 2: Operating Entity Startup	41
<i>Internal Planning and Operations</i>	41
<i>Establishment of Initial NAIRR Resource Components</i>	42
Phase 3: NAIRR Initial Operational Capabilities	44
<i>Initial Computational Resources</i>	44
<i>Initial Data Resources</i>	45
<i>Initial Research Using the NAIRR</i>	45
Phase 4: NAIRR Ongoing (Steady-State) Operations	46
<i>Evolution of NAIRR Resources</i>	46
<i>Partnership Engagement Operations</i>	47
<i>User Outreach, Engagement, and Support Operations</i>	47
<i>Outreach and International Collaboration</i>	47
NAIRR Budget	48
NAIRR Evaluation (Phases 1–4).....	50
Roadmap for Implementation	52
Steps to Initiate the NAIRR in 2023: Actions for the U. S. Government	54
<i>For the President and Executive Branch Departments and Agencies</i>	54
<i>For Congress</i>	54
6. Conclusion	55
Appendix A. Definitions	A-1
Appendix B. Details of NAIRR Task Force Establishment and Approach to Roadmap	
Development	B-1
Charge to the NAIRR Task Force	B-1
Task Force Approach	B-2
<i>Initial Phase</i>	B-2
<i>Final Phase</i>	B-3
Appendix C. Briefers to the Task Force	C-1
Appendix D. Public Input Provided on the Interim Report in Response to the Federal Request for	
Information	D-1
Appendix E. Public Input Provided on the Initial Federal Request for Information on Designing	
the NAIRR	E-1
Appendix F. Subject Matter Experts Consulted by Task Force Members	F-1
Appendix G. NAIRR Public Listening Session.....	G-1
Appendix H. NAIRR Task Force Staff and Contributors.....	H-2
Appendix I. Examples of NAIRR Evaluation Metrics	I-1

Appendix J. Draft Legislative Language for NAIRR Authorization..... J-1
Appendix K. AbbreviationsK-1
Appendix L. Notes L-1

List of Figures

Figure 1. The Current Fabric of U.S. Research Cyberinfrastructure..... 6
 Figure 2. NAIRR Strategic Objective and Goals 8
 Figure 3. A Vision for NAIRR Users and Resource Elements 9
 Figure 4. Proposed NAIRR Governance Structure..... 14
 Figure 5. Process for Selection and Integration of NAIRR Resource Providers..... 16
 Figure 6. Example Elements of a Theory of Change for the NAIRR 50
 Figure 7. NAIRR Implementation Roadmap..... 53

List of Tables

Table 1. NAIRR Six-Year Budget Summary 48
 Table 2. Operating Entity Costs..... 50
 Table 3. Roles in KPI Definition and Frequency of External Evaluation 51
 Table H.1. Examples of Metrics that Could be Associated with a NAIRR Theory of Change I-1
 Table H.2. Examples of Metrics for Assessing Progress toward NAIRR Goals I-2

Task Force Membership

TESS DEBLANC-KNOWLES (CO-CHAIR, BEGINNING AUGUST 2022)

*Senior Policy Advisor, National AI Initiative Office,
White House Office of Science and Technology Policy*

MANISH PARASHAR (CO-CHAIR, BEGINNING OCTOBER 2021)

*Office Director of the Office of Advanced Cyberinfrastructure,
National Science Foundation*

LYNNE PARKER (CO-CHAIR, JULY 2021–AUGUST 2022)

*Former Deputy U.S. Chief Technology Officer and Founding Director, National AI Initiative Office,
White House Office of Science and Technology Policy*

ERWIN GIANCHANDANI (CO-CHAIR, JULY 2021–OCTOBER 2021)

*Assistant Director for Technology, Innovation, and Partnerships,
National Science Foundation*

DANIELA BRAGA

Founder & CEO of Defined.ai

MARK E. DEAN

OREN ETZIONI

CEO, Allen Institute for AI

JULIA LANE

Professor, New York University

FEI-FEI LI

*Sequoia Professor of Computer Science at Stanford University and Denning Co-Director of the Stanford
Institute for Human-Centered AI (HAI)*

ANDREW MOORE

Vice President & General Manager, Google Cloud AI & Industry Solutions

MICHAEL L. NORMAN

Distinguished Professor, University of California, San Diego

DAN STANZIONE

*Executive Director, Texas Advanced Computing Center
Associate Vice President for Research, The University of Texas at Austin*

FREDERICK H. STREITZ

*Chief Computational Scientist, Lawrence Livermore National Laboratory
Senior Advisor, CDC Center for Forecasting and Outbreak Analytics*

ELHAM TABASSI

*Chief of Staff, Information Technology Laboratory,
National Institute of Standards and Technology*

1. Introduction

The economic and national security of the United States has long relied on its unique and vibrant ecosystem for scientific discovery and technological innovation. The United States invests in research and development (R&D) across science and engineering disciplines to advance understanding of natural, built, and human systems and develop tools and methods for solving practical challenges. This R&D leads to downstream development of applications and commercial products that drive economic growth while supporting the human aspiration to explore, understand, and improve the conditions of our world.

AI is increasingly a key driver of the Nation's research and innovation ecosystem, as it holds the potential to power discovery, innovation, and economic growth across every field of science and every sector of the economy. However, achieving this potential and harnessing AI to tackle grand challenges require substantial and sustained investment in AI R&D as well as education and workforce development.¹ It also requires access to the infrastructure necessary for AI experimentation and training. Currently, uneven access to the resources that fuel AI R&D and training have limited opportunities for researchers and contributed to a lack of diversity in the field. This lack of diversity means that the full range of talent is not being leveraged for this work. Lack of diversity may also contribute to the development of biased or harmful AI systems and threaten the Nation's innovation potential and global leadership. Concerns related to misuse of AI and environmental effects of AI development are also increasing. Making computational, data, and training resources available to more of America's researchers through an approach grounded in equity and security can chart a path forward. In this future America can responsibly harness the potential of AI for societal good and economic wellbeing—while also strengthening American technological competitiveness for decades to come.

The Current Landscape of AI R&D

The term "Artificial Intelligence" refers to a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments (see Box 1).²

AI systems can be applied to tasks spanning diverse areas, including planning and optimization, perception and vision, modeling and simulation, natural language understanding, robotic process automation, recommendation, and prediction. These tasks can be accomplished through statistical inference extracted from "training" data (in the case of Machine Learning [ML]) or programmed logical

Box 1. Definition of Artificial Intelligence³

The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to:

- (A) Perceive real and virtual environments.
- (B) Abstract such perceptions into models through analysis in an automated manner.
- (C) Use model inference to formulate options for information or action.

reasoning (as with expert systems). Today, the computational and storage capacity of computer systems has advanced to a stage where storage and analysis of large quantities of data has become not only possible, but also an increasingly dominant enabler of R&D. Parallel development of advanced software tools and algorithms have facilitated realization of powerful analytical and predictive methods based on AI, which are being applied broadly across fields of science and engineering.

AI technologies and sustained investments in cyberinfrastructure have supported scientific and technological breakthroughs in diverse areas such as protein folding, nuclear fusion, and even programming. The breakthroughs did not happen by chance. They emerged from an ecosystem characterized by decades of systematic investments in cyberinfrastructure, education and training, and large and growing amounts of data and computational power; and the rich collaborations between academic researchers and the private sector. The potential for the U.S. research community to contribute to the global AI research and innovation ecosystem is growing.

In recent years, academia has seen a significant growth in AI and computer science research and education. Since 2016, about 2,000 computer science faculty members have published at least one AI-related paper, and on the order of 900 have published at least five.⁴ In 2019 roughly 28,000 undergraduate students received degrees in computer science, more than doubling the number of degrees awarded in 2014. Those who pursue doctoral programs in computer science and related fields in North America are increasingly choosing to specialize in AI: The share of new computer science PhD recipients specializing in AI increased from 19 to 25 percent between 2019 and 2020, for a total of 442 in 2020.⁵

However, increased investments in AI research and education have not been distributed equally across the Nation's researchers and innovators.⁶ Of the U.S. resident AI PhDs conferred in 2020, approximately 51 percent were awarded to non-Hispanic Whites, 30 percent to Asians, 7 percent to Hispanics (compared to their representing 18.9 percent of the U.S. population), and 2 percent to Blacks or African Americans (compared to their representing 13.6 percent of the U.S. population). These numbers show a decrease in the percentage of AI PhDs awarded to Hispanic and Black or African American students relative to 2010.⁷ Similarly, gender diversity in AI is low and has demonstrated little change over the past decade. According to one estimate, about 20 percent of both the AI PhD and computer science PhD graduates in North America in 2020 were female.^{7,8} This lack of diversity among students and graduates gives rise to a corresponding lack of diversity in the workforce, and contributes to the development of AI tools and approaches that perpetuate systemic bias and limits the breadth of ideas incorporated into AI innovation.⁹

While academic and private sector interest in AI has grown, access to the computational and data resources that fuel much of today's AI has become concentrated in large private-sector firms, well-resourced universities, and national laboratories, creating a growing divide that limits innovation and growth.¹⁰ The resulting impact on U.S. innovation and economic growth is evident. Even though private investment in AI more than doubled between 2020 and 2021 to approximately \$93.5 billion, the number of new companies has decreased.⁸ The disparity in availability of AI research resources affects the quality and character of the U.S. AI innovation ecosystem, contributing to a "brain drain" of top AI talent from academic and research institutions to a small

set of well-resourced corporations.¹¹ Such trends have adverse implications for the Nation's capacity to train the breadth of talent required to support future U.S. competitiveness and innovation.

An Opportunity for Strengthening AI R&D in the United States

Sustained investments in AI R&D have enabled the United States to be a longstanding global leader in the field of AI, from the foundations of the field to the present day. Conference papers and AI repository publications by U.S.-based authors remain the most cited globally. However, American dominance is currently threatened. Countries such as China have made long-term investments that are bearing fruit in terms of both their scientific and technological achievements. For example, authors based in China have overtaken U.S.-based authors in AI journal publication citations. The United States has been granted more AI patents than any other nation, although AI patent applications from China far surpass those from the United States.⁸ These trends illustrate the rapidly changing AI innovation landscape as output in AI R&D continues to grow rapidly worldwide and as leadership in AI and other emerging technologies has become a central facet of geopolitical competition.

AI breakthroughs could accelerate progress across a range of mission areas of Federal agencies: from energy and sustainability to healthcare and biomedical treatments to foundational research. For example, AI could support a broad spectrum of actions needed to build a more sustainable future—from mitigation of greenhouse gas emissions and development of data-driven strategies for conservation, to automated solutions for managing consumption and the invention of new clean energy sources and materials.

Realizing the benefits of AI for the Nation will rely on the ability of all U.S. researchers to access the necessary cyberinfrastructure, especially researchers with limited resources or who have been historically excluded from AI and related fields and industries. Engaging the full diversity of U.S. talent will bring important perspectives, research capacity, and inspiring use cases.

Critical opportunities for strengthening the U.S. AI R&D ecosystem exist in four key areas:

- **Innovation:** Bringing together complementary resources, capabilities, and skills could enable new modalities of research, new understanding and knowledge, and new, transformative solutions.
- **Diversity:** Engaging the full breadth of talent in the United States can help introduce new ideas and use contexts for AI, and expand and strengthen the potential of AI R&D in the United States, including for addressing a range of societal challenges.
- **Capacity:** Increasing the development of and access to resources optimized for foundational, use-inspired, and translational AI R&D is essential for supporting a growing AI R&D community and its needs.
- **Trustworthiness:** Practical and societal implications of AI must be considered in all AI R&D, given its potential for ubiquitous application throughout the economy and society. As with any powerful and complex tool, AI comes with risks; responsibility for managing such risks is shared across all phases of the AI life cycle, including R&D.

Supporting research on AI’s societal implications, developing testing and evaluation approaches, improving auditing capabilities, and developing best practices for responsible AI R&D can help improve understanding and yield tools to manage AI risks.

Cultivating a vibrant and inclusive AI innovation ecosystem that reflects American values will drive economic growth, national security, and scientific progress, which will in turn increase America's future technological competitiveness. Such outcomes will not be possible through action by any single sector or entity, but require collaborative action among government, academia, the private sector, and non-profits.¹²

In January 2021, as part of the National Artificial Intelligence Initiative Act of 2020,³ Congress established the National AI Initiative to further coordinate and enhance Federal actions toward four objectives: (1) ensure continued U.S. leadership in AI research and development; (2) lead the world in the development and use of trustworthy AI systems in the public and private sectors; (3) prepare the present and future U.S. workforce for the integration of AI systems across all sectors of the economy and society; and (4) coordinate ongoing AI research, development, and demonstration activities among the civilian agencies, the Department of Defense, and the Intelligence Community to ensure that each informs the work of the others. The Initiative codifies sustained and consistent support for AI R&D through grants, cooperative agreements, testbeds, and access to data and computing resources, and requires that the National AI R&D Strategic Plan that focuses AI R&D investments across agencies be updated every three years.

The National AI Research Resource Task Force

As part of the National AI Initiative, Congress established the National Artificial Intelligence Research Resource (NAIRR) Task Force, calling for it to “investigate the feasibility and advisability of establishing and sustaining a [NAIRR] and to propose a roadmap detailing how [a NAIRR] should be established and sustained.”³ A widely accessible, AI-specific research cyberinfrastructure (as defined in Box 2) could meet the opportunities and challenges described above, in alignment with the National AI R&D Strategic Plan, and help to build a stronger, more inclusive U.S. AI R&D ecosystem. This vision is reinforced by the recent CHIPS and Science Act of 2022, which appropriates funding to accelerate advanced computing development, from next-generation graphics processing units to high-density memory chips, and authorizes investments to help actively engage the full breadth and diversity of U.S. talent in the frontiers of science and engineering, including AI.¹³

Box 2. Definition of NAIRR³

The terms "National Artificial Intelligence Research Resource" and "Resource" refer to a system that provides researchers and students across scientific fields and disciplines with access to computational resources, co-located with publicly available, artificial intelligence-ready government and non-government datasets, and a research environment with appropriate educational tools and user support.

The NAIRR Task Force strongly agrees that a shared, AI-focused federation of cyberinfrastructure resources—including computer hardware, data, algorithms, software, services, networks, and expertise—is necessary to transform the AI R&D landscape in the United States.

More equitable access to computational power, large and unbiased datasets, and software tools is needed to empower a diverse collection of individuals and teams across the country to advance AI methods and technologies; use AI to make progress on science, engineering, and societal challenges; and actively contribute to the development and adoption of AI systems, policies, and practices that respect privacy, civil rights, and civil liberties. The NAIRR Task Force found that developing a NAIRR is both feasible and advisable, and this final report provides the implementation plan to do so.

This report to the President and Congress represents the culmination of the Task Force’s efforts and provides a path forward and specifications for meeting national cyberinfrastructure needs and transforming the AI R&D landscape for the benefit of all. It builds upon and extends the Task Force’s interim report, submitted to the President and Congress in May 2022,¹⁴ which set forth the Task Force’s vision and preliminary recommendations for key aspects of the NAIRR, based on a variety of information-gathering and public input, as indicated in Box 3. (See Appendix B for more details on how the Task Force conducted its work. Appendix F lists subject matter experts who briefed the Task Force, while Appendix G provides information on the public listening session.)

To succeed, the NAIRR must be designed to leverage and complement the existing cyberinfrastructure fabric for R&D across the Nation—and augment or supplement it accordingly. The current fabric spans high-performance and leadership computing facilities, distributed computing frameworks, commercial cloud resources, and the networks that bring them to users; data; software and tools; testbeds; and educational tools and programs (see Figure 1). A successful NAIRR must also foster the participation of individuals and groups across sectors and domains in AI R&D, and provide opportunities to include the expertise and experience of all stakeholders.

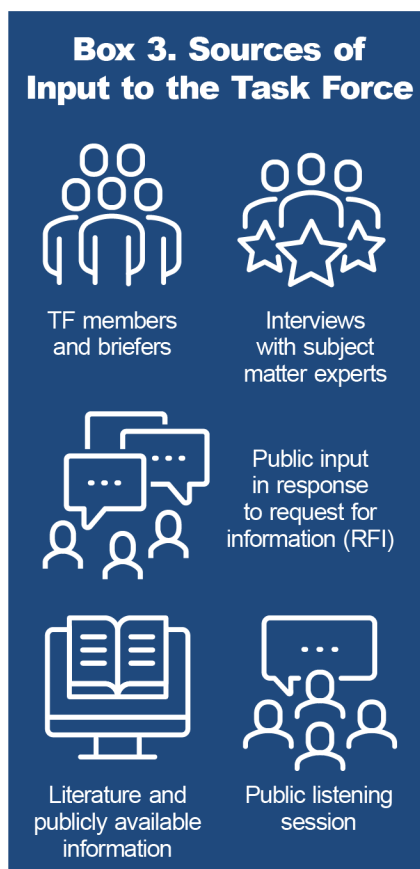




Figure 1. The Current Fabric of U.S. Research Cyberinfrastructure

Structure of This Report

The following chapters set out a roadmap for the NAIRR, including key implementation steps, attributes, and specifications necessary for the NAIRR to fulfill its purpose. Chapter 2 describes the vision and goals for the NAIRR and identifies responsible entities and a general timeline for its establishment. Chapter 3 describes key attributes for NAIRR governance, technical resource components, security and user access controls, and user training and education tools and strategies. Specific actors and actions are identified to the extent possible at this time, recognizing that many decisions will wait until implementation or be revisited then. Chapter 4 provides more detailed specifications for NAIRR resource components, including NAIRR initial operational capabilities (i.e., the set of initial resources and functions that must be in place to launch NAIRR operations). Chapter 5 describes a phased buildout plan for establishing NAIRR governance, management, resources, and operations. Chapter 5 also provides a budget estimate for establishment and sustainment of the NAIRR, a list of actions for each buildout phase, and immediate next steps for U.S. Executive Branch agencies and Congress.

2. A National Cyberinfrastructure to Democratize and Accelerate AI R&D

NAIRR Vision and Goals

The NAIRR is envisioned as a widely-accessible, national cyberinfrastructure that will advance and accelerate the U.S. AI R&D environment and fuel AI discovery and innovation in the United States by empowering a diverse set of users across a range of fields through access to computational, data, and training resources. Created by leveraging, linking, and augmenting the Nation's existing cyberinfrastructure resources, the NAIRR would support cutting-edge explorations in AI R&D and improve the ease of collaboration across disciplines and sectors that address pressing problems with AI. It would create opportunities to train the future AI workforce, support and advance trustworthy and responsible AI, and catalyze development of ideas that can be practically deployed for societal and economic benefits.

The NAIRR would accelerate these outcomes by enabling U.S.-based researchers to access the digital resources that enable AI R&D: computational power, datasets, software tools, and training and collaboration resources. These would be made available through an integrated user portal with key user functionalities such as single sign-on access to resources, collaboration tools, search tools for resource discovery, detailed resource specifications and user guides, an interface for computational job submission, and consolidated accounting of resource use. Researchers would be able to request computational allocations across a range of high-performance computing (HPC), commercial cloud, and other remote, on-premises or distributed computing resources. User support services and interactive training modules would support users new to the field, which, along with clearly-defined policies and standards of practice, would promulgate best practices for trustworthy AI model development and responsible data use by design. A publicly-accessible NAIRR user portal would provide curated catalogs that list commonly-used AI datasets, testbeds, educational resources, and relevant metadata, serving as a clearinghouse for the AI R&D community. Through a tiered-access model, vetted researchers would be able to conduct research on sensitive or restricted data in secure enclaves.

The Task Force believes that the objective for establishing the NAIRR should be to strengthen and democratize the U.S. AI innovation ecosystem in a way that protects privacy, civil rights, and civil liberties. The NAIRR objective will be achieved by pursuit of four measurable goals: (1) spur innovation, (2) increase diversity of talent, (3) improve capacity, and (4) advance trustworthy AI, as illustrated in Figure 2.



Figure 2. NAIRR Strategic Objective and Goals

The NAIRR User Base

The NAIRR should support the needs of researchers and students from diverse backgrounds who are pursuing foundational, use-inspired, and translational AI research. The users of the NAIRR are envisioned to fall into three primary categories:

- (1) **Researchers conducting AI research:** those who advance the state of the art in AI or understanding of its sociotechnical dimensions, or those who develop innovative applications of AI to solve problems in another domain of study (while also furthering AI itself), including science, engineering, medicine, business, education, and the humanities.
- (2) **Educators incorporating AI tools and training resources into learning environments:** for example, through classroom demonstrations, homework assignments, and interactive experiences.
- (3) **Students learning about AI:** those studying at community colleges, four-year colleges and universities, or graduate schools who are learning and experimenting with the development of AI models, tools, and applications as well as exploring the societal and economic implications of AI innovations; and those pursuing re-skilling programs in AI.

The primary user groups of the NAIRR should be U.S.-based and affiliated with U.S. academic institutions; non-profit organizations; Federal agencies or federally funded research and development centers (FFRDCs); State, local, or Tribal agencies; and startups or small businesses that have been awarded Federal grants via the Small Business Innovation Research (SBIR) or Small Business Technology Transfer (STTR) programs, or other similar Federal programs, for small businesses to advance foundational, use-inspired, or translational AI R&D.

To ensure that the AI research ecosystem is diverse, the NAIRR should aim to transform its users’ capabilities and outcomes. For example, the NAIRR should reduce barriers to participation in AI R&D and education, and make it easier and less costly for researchers—especially those who have not historically been engaged and have been underrepresented in AI R&D—to access key AI research tools. To ensure that there is ample workforce capacity, educators should have new, readily available options for incorporating AI tools and training materials that support student learning in AI, including the ethics of AI. Students should gain new and early exposure to AI tools and methodologies that transform their understanding; increase their interest in AI and other science, technology, engineering, and mathematics (STEM) fields; and broaden engagement across the full pool of talent to help build a strong and diverse future AI innovation ecosystem. A vision for how users will access and benefit from the NAIRR is illustrated in Figure 3. To maximize the impact of the NAIRR, complementary agency programs could also be initiated, with associated Federal appropriations, to support the entry of new researchers into AI R&D who may then leverage the NAIRR, as a parallel means of growing, diversifying, and democratizing the R&D community.

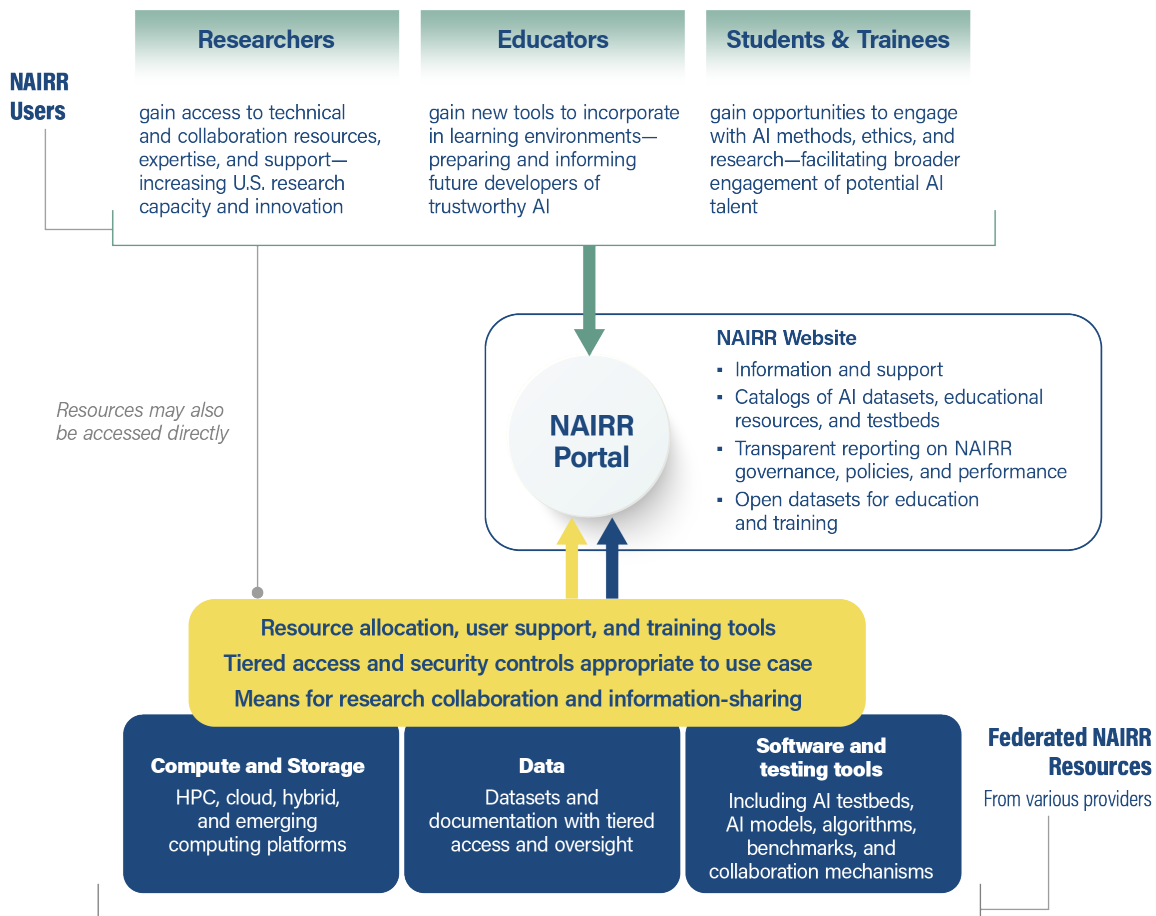


Figure 3. A Vision for NAIRR Users and Resource Elements

NAIRR Constituents

The success of the NAIRR will hinge on the leadership, participation, and engagement of a diverse mix of organizations, groups, and researchers across a range of sectors and disciplines. Government, academia, industry, and civil society groups will all have critical roles to play in realizing the vision of the NAIRR.

Government

The U.S. Government should have the primary role in establishing the NAIRR. The Federal Government should be its principal sponsor, funding NAIRR to help meet its goals in the national interest and the government-wide National AI Initiative, which involves activities across Federal agencies and is coordinated by the National AI Initiative Office (NAIIO) within the White House Office of Science and Technology Policy (OSTP). Federal departments, agencies, and offices that conduct or support AI R&D or provide research cyberinfrastructure should take active roles in supporting the establishment and governance of the NAIRR and funding its component resources.

The government has a strong foundation on which to build the NAIRR. Many Federal agencies already support AI R&D and R&D cyberinfrastructure. As reported in the Networking and Information Technology R&D Program and the NAIIO Supplement to the President’s fiscal year (FY) 2023 Budget,¹⁵ 11 Federal departments plus the independent agencies National Science Foundation (NSF) and the National Aeronautics and Space Administration (NASA) reported investments in AI R&D. Four of these departments and agencies (the Department of Defense, the Department of Health and Human Services, the Department of Energy [DOE], and NSF) reported funding more than \$200 million dollars each in AI R&D in FY 2022.

Many Federal agencies are making important strides in using AI to advance their agency missions—from improving education outcomes to transforming the detection and treatment of diseases (and much more). Their work could be accelerated by research facilitated through the NAIRR. As a national resource, the NAIRR could be leveraged by agency researchers and supported by agencies through the multi-agency governance structure described in Chapter 3.

Federal agencies (including via their FFRDCs) can also contribute research resources to the NAIRR, such as large datasets, computing resources, software tools, and AI testbeds. State, territorial, local, and Tribal governments may also contribute datasets suitable for research, and could benefit from the results and applications of research performed through the NAIRR.

Academia

The NAIRR should provide researchers, educators, and students at universities and colleges across the United States with access to the computational and data resources that fuel cutting-edge AI research, along with training materials and user support. The NAIRR offers particular value to institutions whose researchers have not historically received significant Federal AI research funding or cyberinfrastructure support, or whose lack of resources has inhibited participation in the AI R&D enterprise. The NAIRR thus offers opportunities to broaden participation in AI research, complementing provisions in the CHIPS and Science Act of 2022 aimed at strengthening research capacity and expanding STEM education opportunities in emerging technologies at

historically Black colleges and universities and minority-serving institutions such as Tribal colleges or universities and Hispanic-serving institutions. Its accessible education resource catalogs and training tools will offer value for learners from diverse backgrounds, organizational affiliations, and geographic locations.

It will be critical to ensure that universities and their researchers have an important role in establishing and managing the NAIRR for several reasons. First, academic researchers and research groups will be vital users of and contributors to the collaborative resources such as datasets and research tools available through the NAIRR. Second, academic researchers engaged in cutting-edge research will be key to providing strategic advice and oversight for the NAIRR's investments in computational and data resources. Finally, universities are the front line in designing the curricula and training materials that are necessary to expand the capacity of a diverse AI workforce.

Industry

Companies should benefit from the flow of a diverse group of graduates whose training is supported by NAIRR resources and from the innovations resulting from NAIRR-supported research. Startups and small businesses should have the opportunity to use NAIRR resources for their own R&D.

For-profit and not-for-profit organizations have products and services that could be made available through the NAIRR, and thus should also have the opportunity to provide resources for inclusion in the federated cyberinfrastructure—potentially through commercial cloud computing contracts or through the incentivized contribution of software tools or datasets. As the NAIRR evolves there should be opportunities for companies to provide funding or other contributions towards the NAIRR's operations through partnership agreements. Industry experts may also participate as technical advisers on NAIRR advisory boards.

Civil Society

The NAIRR should be a platform on which researchers can study and examine societal implications of AI and to develop and test solutions that would maximize the benefits of AI. A variety of scientific and advocacy groups—scientific societies and associations; groups concerned with data privacy, civil rights, and civil liberties implications of AI; philanthropic organizations; and academic researchers—should have the opportunity to leverage the NAIRR for research and evaluation that promote the responsible development and use of AI. Scientific and advocacy groups could also participate in oversight of the NAIRR as members of advisory boards. They should play an important role in ensuring that public interests, such as the development of trustworthy AI, are properly represented and considered among NAIRR governance and management entities.

3. NAIRR Organization, Management, and Governance

The impact of AI extends to nearly all sectors of the Nation’s economy and aspects of society. Thus, it is critical that the governance of the NAIRR appropriately reflects not only the breadth and diversity of the users of the NAIRR, but also the broad suite of constituents likely to be impacted by the AI innovations that result.

The organizational structure for NAIRR management and governance should incorporate the interests and perspectives of the many Federal agencies involved in AI R&D, take advantage of the distributed nature of existing and future cyberinfrastructure, and facilitate input from the various constituents and communities involved in and affected by AI research. This chapter lays out the recommended organizational structure and management elements of the NAIRR. It also describes the key governance functions that will require policies and oversight, such as building considerations of privacy, civil rights, and civil liberties into all facets of the NAIRR’s design and operations as well as ensuring system security.

NAIRR Organizational Structure

Federal agencies currently invest in the infrastructure that enables federally funded research via a range of different models, in alignment with their mission needs. While management of the NAIRR could be handled entirely within a single government agency (which has the benefit of clear ownership, authority, and responsibility), excluding other agencies would likely narrow its focus to that agency’s specific mission, leaving the needs of researchers supported by other agencies unmet, and translating to a loss of opportunity for the Nation.

Instead, the Task Force recommends that one agency serve as the “administrative home” for the NAIRR to provide core funding for a third-party (non-government) Operating Entity that carries out the activities needed to coordinate, federate, and sustain the NAIRR. This funding would provide for the operations of the Operating Entity, not the research resources that would be a part of the NAIRR. Other agencies should play a major role in NAIRR stewardship by (1) forming a multi-agency Steering Committee that provides strategic guidance and collective oversight of the NAIRR, (2) funding resource providers that would be federated together to constitute the NAIRR, and (3) providing staffing for the Program Management Office. It is critically important that all agencies involved in the NAIRR work together through the Steering Committee to coordinate the provisioning of resources and ensure that all agency perspectives are reflected. The Task Force majority recommendation for the NAIRR administrative home is described in Box 4.

The Task Force makes its recommendations after careful review of the successful cooperative stewardship model for multidisciplinary users of the Nation’s synchrotron, neutron, and high-magnetic-field user facilities reviewed by the National Research Council.¹⁶ In this model, the responsibility for design, construction, operation, maintenance, and upgrading of a research facility core rests with a single clearly identified Federal agency—the steward. The steward then engages partners—other agencies, industry, and private institutions—in the planning, design, construction, support, and funding of the experimental stations and other sub facilities. While no model is without flaws, the Task Force believes this model will best serve the AI R&D priorities across Federal agencies and achieve the societal-level impacts envisioned for the NAIRR.

Leveraging this model, the agency serving as the administrative home for the NAIRR would fund and oversee the core operations, but would not establish the strategic direction of the NAIRR, nor fund all the individual resource providers. As described below, a Steering Committee, with representation from agencies participating in the NAIRR, should set the strategic direction of the NAIRR and drive decisions about which resources will be brought into the NAIRR from which providers.

Given the complexity of the NAIRR, the Task Force recommends that its day-to-day operations be managed by a single, non-governmental Operating Entity. The Operating Entity will require a dedicated, expert, stable workforce composed of highly trained technical talent capable of managing long-term, complex needs and systems with a high degree of objectivity. The Operating Entity must leverage external input-gathering mechanisms. Given the NAIRR’s many operational requirements, expert advice is needed on issues spanning technical resource design, development, management, interoperability, standards, and improvement; user experience design, development, and improvement; ethical design, development, and use of research resources; legal and regulatory compliance, intellectual property management and agreements; and education and training. Experts from a wide range of scientific and academic disciplines, including social science and ethics, and also drawn from government, industry, and non-profit sectors, must therefore be actively engaged, for example, by including them among members of the advisory boards and a

Box 4. Designating the NAIRR Administrative Home

The Task Force, by majority opinion, recommends the designation of NSF as the administrative home for the NAIRR. The Task Force defined the key attributes envisioned for an effective administrative home to include the following:

- Mission alignment.
- Capacity and capabilities to effectively support administrative activities.
- Existing relationship with the AI research community and other NAIRR constituents.
- Experience supporting foundational, use-inspired, and translational AI research.
- Existing relationship to building workforce capacity at multiple levels.
- Focus on equity and diversity and the ability to support democratization of resource access.

The Task Force assessed that NSF meets these attributes and could effectively oversee the operations of the NAIRR within the collaborative interagency framework proposed. NSF’s relationship with America’s research community in the field of computer science and across all domains of science and engineering, as well as its experience in funding broadly-used national cyberinfrastructure resources, services, and expertise, provides it with the existing relationships, trust, and expertise necessary for a rapid and effective stand up of the NAIRR.

User Committee. These advisory bodies are intended to bring diverse perspectives, providing strategic management advice to inform the NAIRR’s operations.

The recommended organizational structure for the NAIRR (see Figure 4) shows how different elements of the NAIRR management and governance structure should relate and interact. A detailed description of each of these elements is provided in the sections that follow.

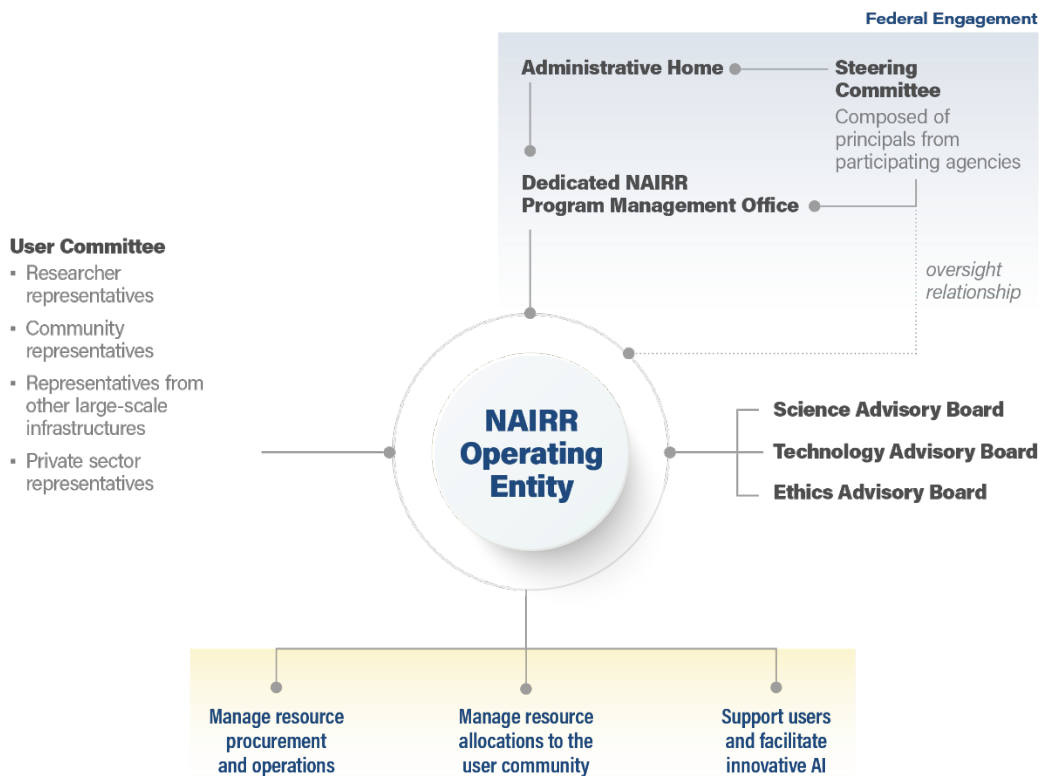


Figure 4. Proposed NAIRR Governance Structure

Steering Committee

Many Federal agencies individually and collectively have stakes in the NAIRR’s success, and are therefore envisioned to contribute to its governance. NAIRR governance should follow the proposed cooperative stewardship model and serve the interests of all agencies involved. A Steering Committee comprising principals (e.g., deputy or assistant secretaries) at departments, agencies, and offices with significant AI R&D investments or equities in the NAIRR should be constituted to provide strategic direction. This Steering Committee should be chaired by the Director of the NAIIO, in accordance with the office’s role as coordinator of Federal activities in support of the National AI Initiative, and should have rotating co-chairs. The involvement of deputy or assistant secretaries ensures top-level commitment to agencies’ engagement in the NAIRR and its governance. The Steering Committee may establish operational working committees to manage more operational issues.

Agencies that have already made substantial investments in AI R&D and cyberinfrastructure are likely most able to provide guidance about the NAIRR’s initial setups and structure, and

therefore are most likely equipped to lead the initial phases of the NAIRR's development. Since all agencies stand to benefit, additional agencies should be brought into the Steering Committee over time. The Steering Committee composition should be reviewed every three years by the NAIIO. As part of these periodic reviews, additional agencies could commit funding or resources to the NAIRR or become members of the Steering Committee, or participating agencies could elect to discontinue participation.

The Steering Committee will establish the overall strategic direction for the NAIRR and should be responsible for overseeing and approving the following:

- The operating plan, budget (see Chapter 5), and requests for proposals (RFP) to solicit bids for the Operating Entity, including the terms and conditions and functions of the Operating Entity.
- The review of proposals for and selection of the awardee to serve as the Operating Entity.
- The identification of resources that could be federated, selection of individual resource providers, and determination of how resources could be allocated and made accessible via the NAIRR.
- Once the NAIRR has been initiated, the development of key performance indicators (KPIs) for the Operating Entity and NAIRR as a whole, in collaboration with other NAIRR governance entities.
- Work with an external, independent evaluator to conduct a periodic review of NAIRR activities and performance against KPIs, and assess program needs and inform decision making and planning.

The Steering Committee should initiate work on the above areas through the administrative home and NAIRR Program Management Office, and may manage certain of the above tasks through operational working committees. The Steering Committee should monitor the progress of the NAIRR towards its objectives and provide recommendations annually in a publicly available report to the NAIIO.

Individual Agencies

Federal agencies with AI R&D investments or equities should contribute NAIRR resource elements by incorporating appropriate funding for NAIRR resources in their annual budget requests. Funding for core operations of the NAIRR through the Operating Entity should be provided by the agency serving as the administrative home; individual resource elements can be funded separately with provisions for federation as part of the NAIRR.

First, funding should be directly allocated by Congress to the agency serving as the administrative home for the NAIRR to provide for the activities of the Operating Entity, including project management, portal development and deployment, federation support, and concierge services such as training and user support. The administrative home agency should staff a Program Management Office, which is described in detail in the next section.

Second, funding should be directly allocated by Congress to individual agencies to fund the resources made available through the NAIRR, many of which may be aligned to specific agency mission interests, which should be federated together to constitute the NAIRR. Resources can be funded individually or as part of multi-agency funding opportunities coordinated through the Program Management Office. In addition to software and platform-as-a-service providers, the NAIRR resource providers may represent one or both of the following:

- Expansions of existing computing capacity (e.g., on-premise computers at a university center or at an FFDRC), dedicated computing time or storage purchased from commercial cloud computing providers, or purchases of new, specialized computational facilities dedicated to AI research.
- Trusted data providers and hosts for a transparent and responsible AI data commons. Access to data should be tiered, controlled by the data providers, and provided through the same portal through which computational resources are provided.

Given the costs of these resources and their broad applicability to many types of AI R&D and research using AI-enabled methods, in some cases it will be more efficient for multiple agencies to collaborate in funding NAIRR resources rather than having each participating NAIRR agency individually purchase and contribute computing and data storage resources to the NAIRR. Additional context about the process for selecting and integrating resource providers into the NAIRR is provided in Figure 5.

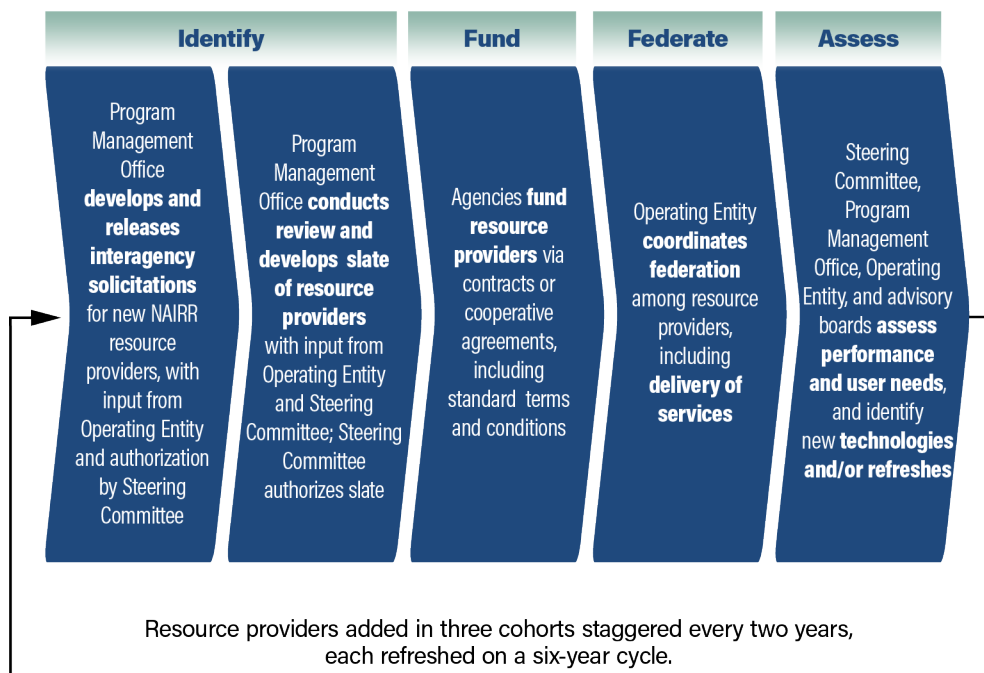


Figure 5. Process for Selection and Integration of NAIRR Resource Providers

Third, appropriations provided to Federal agencies for AI R&D programs should be sufficient to support inclusion of NAIRR allocations to enable access to AI research resources as part of Federal awards to investigators funded through agencies' own intramural and extramural proposal

and review processes. Including NAIRR resources as part of the awards will enable such federally funded researchers to leverage the NAIRR's full capabilities in meeting agencies' R&D objectives.

Program Management Office

While the Steering Committee should provide strategic direction for the NAIRR and have ultimate accountability for its success, the scale and complexity of the NAIRR would require ongoing operational oversight and management by Federal Government employees through a dedicated NAIRR Program Management Office. The Program Management Office should include 8 to 10 dedicated Federal agency staff members, including experts in cyberinfrastructure, data, AI R&D, scientific integrity, ethics, and other areas necessary to execute the Steering Committee's vision; staffing could be expanded as needed. The Task Force recommends that the Program Management Office staff include individuals who are on detail from participating agencies, including for leadership positions in the office. In practice, the Program Management Office should serve as the operational arm of the Steering Committee and do the following:

- In consultation with the Steering Committee, develop the solicitation and solicit bids for the Operating Entity, which includes the identification of key Operating Entity personnel such as the Director and key staff.
- Manage the review process and recommend an award by the administrative home agency for the funding of the Operating Entity.
- Identify an external independent evaluation entity whose independent assessment would inform periodic review of the Operating Entity and the NAIRR by the Steering Committee and Program Management Office.
- In collaboration with the Operating Entity, develop multi-agency funding opportunities for resource providers.
- In collaboration with the Operating Entity, manage the review of responses to multi-agency resource provider funding opportunities.
- Administer the Operating Entity contract (i.e., oversee operations/processes including federation of resource providers, assess the Operating Entity's performance on a recurring basis).
- Oversee resource allocation and utilization.

Selection of the Operating Entity should be conducted in consultation with the Steering Committee and through a standard solicitation process. Criteria to guide the selection process should be developed by the Steering Committee, and should include but not be limited to experience managing multi-agency initiatives; identification of key staff personnel; expertise in AI R&D; a strong diversity plan; and an ability to execute according to the NAIRR implementation timeline presented in Chapter 5.

Operating Entity

The Operating Entity should be a distinct, non-government organization, governed by a formal charter and associated policies, with an executive leadership team managing day-to-day operations. It may take the form of an independent legal entity or a consortium of one or more partners (e.g., existing organizations such as research universities, industry, laboratories, and FFRDCs) that work jointly to initiate, manage, and sustain the NAIRR. The Operating Entity should not itself operate the totality of the computer hardware that makes up the NAIRR; instead, computing, data, and training resources would be delivered by resource providers at universities, FFRDCs, and from the private sector. The Operating Entity would manage the day-to-day operations of the NAIRR. It would have five major responsibilities: (1) linking and coordinating the provisioning of federated NAIRR resources; (2) developing NAIRR policies and procedures; (3) continually modernizing the NAIRR; (4) advancing diversity, equity, inclusion, and accessibility (DEIA) in all aspects of the NAIRR, including operations; and (5) establishing mechanisms to enable evaluation, oversight, and the collection of data for assessing KPIs. These responsibilities are described further below.

Coordinate the Provisioning of NAIRR Resources

The Operating Entity should work with the Program Management Office (with guidance from the Steering Committee) to develop one or more multi-agency funding opportunities for resource providers. While agencies may opt to fund resource providers separately, a multi-agency funding opportunity would optimize federation and coordination of individual resource providers. The Steering Committee or their designees should review proposal submissions (in concert with the Program Management Office and Operating Entity) and select awardees. From awards made through the multi-agency funding opportunity process, agencies would contract for resource providers to provide services to the NAIRR, using contracts based on a set of common terms and conditions. In some cases, an agency might provide funding to the Operating Entity for direct contracting of services, such as from cloud providers. Subsequently, the Operating Entity will provide continuous management oversight and service delivery evaluation of resource providers in the context of their federation within the NAIRR, including creating the ground rules for interoperability across resource providers. The Operating Entity will be responsible for working with the providers to implement course corrections as needed. It will also receive and evaluate, on a yearly basis at a minimum, requests from the User Committee regarding what resources the NAIRR should offer.

Develop and Communicate NAIRR Policies and Procedures

The Operating Entity must transparently communicate which individuals or groups are eligible to use the resources, how resources will be allocated among interested users, and how the users will be able to request and gain access to the resources. Thus, the Operating Entity, in consultation with the NAIRR advisory boards and the Steering Committee, will need to establish the corresponding policies and procedures. As part of this effort, the Operating Entity must establish review processes grounded in principles of scientific integrity and ethics to allocate resources fairly, equitably, and transparently for the full diversity of users and user types, including those who have long been underrepresented in AI R&D. To support these efforts, the Operating

Entity will develop portals and services with information about how to access and use resources; hire personnel to serve as the central support staff for NAIRR users and to produce documentation on its use; and create open funding opportunities and associated review processes for project proposals to use the NAIRR's computing resources. When possible, the Operating Entity should leverage existing approaches, such as review processes, employed by Federal research funding agencies.

Continually Update the NAIRR with the Latest Technologies and Capabilities

The Operating Entity should manage a continual updating of the NAIRR infrastructure to include the latest computational, networking, and data collection, storage, and dissemination technologies and capabilities through biennial multi-agency funding opportunities. In collaboration with the User Committee and informed by metrics related to NAIRR resource usage and KPIs for the NAIRR more generally, the Operating Entity should regularly identify new areas for innovation and investment and their requirements from a NAIRR perspective, and work with the Steering Committee to scope the biennial funding opportunities accordingly. This ongoing refresh of resources is critical for the NAIRR to be able to power AI R&D at the cutting edge. The Operating Entity should have primary operational responsibility for vetting resources that become part of the NAIRR, including recommending to the Steering Committee when to onboard and sunset individual resources, and authority to set the standards for the security configurations of these resources. As an independent organization, the Operating Entity will have flexibility in contracting, partnering, or entering into other agreements with individual resource providers, with oversight provided by the Program Management Office and the Steering Committee. NAIRR operational responsibilities will be distributed among the Operating Entity, federated resource providers, and possibly contractors via partnerships or other agreement types, depending on the Operating Entity's needs.

The Operating Entity should provide annual reports, including the contributions of resource providers, to the Program Management Office and the Steering Committee, and make these publicly available. To be fully transparent and accountable about how and why individual resources or resource providers are selected or no longer supported, reports will include a set of recommendations to the Steering Committee regarding how to augment, reallocate, or reduce the NAIRR's offerings.

It is likely that needs will emerge that must be addressed in a timely manner. Another mechanism for identifying emerging needs related to the NAIRR's infrastructure will be for the Operating Entity to conduct a range of activities (with guidance from the User Committee) to solicit input from scientific and user communities and agencies, such as through investigator-initiated workshops to scope emerging areas of science and technology. In addition, the Operating Entity should maintain awareness of computational, data, training, and other infrastructure advances, and strive to make these cutting-edge developments available to the community either through contracts with resource providers executed through the multi-agency funding opportunities or through internal discretionary development funds (e.g., on an initial pilot basis).

Advance Diversity, Equity, Inclusion, and Accessibility

The Operating Entity must be explicitly responsible for incorporating DEIA into all aspects of the NAIRR, including the AI R&D that the NAIRR enables. A DEIA focus should be built into the overall organization, operational plan, and federated system of resources from the beginning, rather than as an afterthought. Extending access to AI research resources as broadly as possible, and incorporating a diverse set of viewpoints into the prioritization of investments, the review of resources and resource providers, and the evolution of the AI research ecosystem, are core to the NAIRR's diversity and capacity goals. NAIRR user access policies therefore must be grounded in the principles of equity, fairness, and accessibility. Assessment of progress and input on engagement with and support of a broad and diverse AI community will be a key aspect of NAIRR governance and oversight activities.

Establish Data Collection, Evaluation, Governance, and Operational Oversight Mechanisms

The Operating Entity should establish mechanisms for monitoring system and organizational performance, including by designing appropriate metrics-collection mechanisms into the system architecture. It will need to engage with an independent, external evaluator to support the review, and create a process for updating organizational and operational procedures as issues are identified. As part of its key role in NAIRR governance, the Operating Entity will also need to define ethics and scientific integrity policies, as well as mechanisms for reporting, adjudicating, and remediating any violations, with guidance from its advisory boards and the Program Management Office.

NAIRR Staff and Executive Leadership Team

The Operating Entity should have an executive leadership team—including a Director, Chief Executive Officer, and Chief Operating Officer—that is responsible and accountable for day-to-day operational decision-making for NAIRR operations; interfacing with advisory groups and government oversight entities; managing outreach, communications, and partner engagement; and scouting and strategizing for new and emerging AI R&D needs.¹⁷ Importantly, the Operating Entity Director or executive leadership team should be allocated 5–10 percent of total resources for discretionary allocations; these allocations could be leveraged during emergency situations, allowing the NAIRR to be agile in responding to urgent or atypical needs—for example, as was done with research efforts established in response to the emergence of the COVID-19 pandemic.

To support its responsibilities and functions, the Operating Entity must be able to hire and retain high-quality and experienced staff. For example, ensuring that the NAIRR is resourced with cutting-edge technologies and capabilities requires that the Operating Entity comprise staff members who are expert in advanced research cyberinfrastructure. Similarly, promoting equitable access to resources requires that the Operating Entity's leadership understands barriers to access. The Operating Entity will need to explore a range of mechanisms for making the work of the Operating Entity attractive to an expert, dedicated staff. In addition, for the NAIRR to successfully promote diversity, equity, and inclusion in AI, it must embody these principles by ensuring diversity among its own staff and leadership and enlisting experts with a range of backgrounds and experiences.

NAIRR Advisory Boards

Since the NAIRR will serve many communities and have so many operational requirements, the Operating Entity will need advice on a variety of operational issues, including (1) technical resource design, development, management, interoperability, standards, and improvement to ensure that the NAIRR remains at the cutting edge of innovation; (2) user experience design, development, and improvement to ensure broadly available and equitable access and use of research resources; (3) ethical design and development of access protocols and mechanisms; (4) legal and regulatory compliance; (5) intellectual property management and agreements to ensure that the NAIRR is—and is seen as—trustworthy; and (6) education and training to meet the workforce capacity needs of the AI ecosystem.

To ensure that the NAIRR meets its objective and goals, the Operating Entity should establish several independent boards, focused on different aspects of the NAIRR's mission (e.g., science and technology, data policies, ethics, privacy, civil rights, and civil liberties). These oversight boards and advisory boards should be tasked with providing guidance in specific areas and input on metrics to be used for evaluation.

To this end, the NAIRR should establish at least four advisory boards:

- A **Science Advisory Board** to provide advice about the rapidly changing needs across multiple scientific domains so that the NAIRR can rapidly adapt to support innovation. The Science Advisory Board should include individuals with management experience drawn from the scientific community, the public at large, public interest groups, the private sector, and other large-scale cyberinfrastructure projects.
- A **Technology Advisory Board** to advise the Operating Entity about cutting-edge technological solutions in the provisioning and use of computational and data infrastructures, workforce training, and on privacy- and security-related technologies. The Technology Advisory Board should include recognized experts from across the computing, data, and security communities and should be selected to represent industry and government, with some academic involvement.
- An **Ethics Advisory Board** to advise the Operating Entity on issues of ethics, fairness, bias, accessibility, and AI risks and blind spots. The Ethics Advisory Board's intended roles are to (1) evaluate the ethical use of AI, computational, and data resources by NAIRR awardees as well as issues related to scientific integrity, and help the Operating Entity ensure that privacy, civil rights, and civil liberties are not violated; (2) evaluate and advise on the fairness and appropriateness of data and training delivered by the NAIRR; (3) provide guidance on approaches to understanding issues of ethics, bias, and fairness and on NAIRR ethics policies and practices; and (4) handle concerns and/or complaints brought to the Operating Entity's attention or by the User Committee. The Ethics Advisory Board should provide periodic insight and feedback on a broad range of policy issues, guidelines, and practices, including in areas such as privacy, civil rights, and civil liberties. The Ethics Advisory Board should be selected to include

experts in privacy, civil rights, civil liberties, and ethics as well as to represent user groups, scientific societies, advocacy and civil society groups, and government.

- A **User Committee** to provide the user perspective for the NAIRR, providing feedback on operational and governance issues, offering perspectives on user needs and requirements, and identifying new directions for the NAIRR to create value and serve the community. It should be composed of subject matter experts from across multiple scientific and user communities and be selected to represent AI researchers, with some industry and ex-officio government representation.

The activities of these advisory boards should be supported by staff at the Operating Entity. As a guiding principle, each board should consist of 6–8 members to be selected by the Operating Entity after an open call and with input from the Program Management Office and Steering Committee. Special attention should be paid to diversity, inclusivity, and representation/affiliation of board membership. The exact number of, and nominees for, these advisory boards should be reviewed on a regular basis by the Operating Entity in consultation with the Program Management Office as the number of domains supported by and types of services provided by the NAIRR evolve. Members should represent government, academia, and industry sectors, with the relative weights appropriate for each board. Care should be taken to address potential conflicts of interest. The term of membership for individuals should be three years, with staggered expirations (e.g., one-third rolling off each year). The members of each board will select a chair from among their ranks, who can serve an additional two years in this capacity. Advisory boards report to the Operating Entity executive management and are responsible for delivering written guidance annually. Board reports will be shared with the Program Management Office and the Steering Committee by the Operating Entity. Each board should meet a minimum of twice a year.

Evaluation Entity

Evaluation of NAIRR performance—toward both its high-level goals and its operational KPIs—should be conducted by an independent, external evaluator with experience in assessing major R&D infrastructure programs. This entity should be contracted by the Program Management Office with the input from the Steering Committee, and its evaluation approach developed in parallel with Operating Entity activities so that appropriate metrics can be developed and the associated data collection may be incorporated into the NAIRR’s design.

User Access and Resource Allocation

Since the fundamental objective of the NAIRR is to democratize access to AI resources, the NAIRR must primarily be sustained through Federal investment, with direct user fees employed only to scale beyond a base level of resources. As described in Chapter 2, the primary users of the NAIRR would be U.S.-based AI researchers and students at U.S. academic institutions, non-profit organizations, Federal agencies or FFRDCs, or startups and small businesses awarded SBIR or STTR funding. Others (e.g., private sector researchers other than small businesses) would be allowed to access NAIRR resources, but only at limited levels and in support of research that is in the public interest. Supporting the academic research community should be prioritized through the resource allocation process, with particular attention to underserved communities.

Access to the NAIRR should be granted directly to researchers by Federal funding agencies or the NAIRR Operating Entity. Awards may be flexibly structured to include in-kind credits or tokens redeemable for computer time, data access, or other services.

With oversight and approval from the Steering Committee and Program Management Office, the Operating Entity should establish multiple allocation processes based on the nature, size, and scope of the requests, which are divided along two primary tracks: one driven by participating agencies and a second peer-review track run by the Operating Entity. Within the agency-driven track, agencies should be given latitude in how to make awards, within the constraints of their allocated credits and in close coordination with the Operating Entity. Credits could be awarded directly through agency research grant funding programs or could be made to awardees through a separate process managed by the agency in close coordination with the Operating Entity. Because not every participating agency may have the expertise or resources to run such a process, the agency could choose to leverage the peer-review track managed by the Operating Entity. The Operating Entity should be responsible for keeping the agencies within their allocation caps, which would be determined based on a combination of factors such as an agency's support of AI R&D, contributions of resources to NAIRR, or number of allocation requests received, while enabling the agencies to decide who receives the allocations.

The peer-review track should be managed by the Operating Entity and subdivided by size and type as follows:

- **Startup requests:** These requests should be capped at a modest size (e.g., suitable for a classroom of students for a single semester, or approximately \$1,000 worth of computational time/storage). Requests should be reviewed by staff at the Operating Entity, with turnaround times to the applicant of less than two weeks. Startup allocations would typically expire in one year and then could be renewed.
- **Research requests:** Larger requests in support of significant AI research projects should be peer reviewed through the Operating Entity. The Operating Entity should organize review panels quarterly, and should place caps on the size and duration of requests based on the capacity of resources within the NAIRR.
- **Purchases:** Users could opt to purchase additional allocations if they need services that extend beyond the amount they can acquire through the open startup and research tiers, or could be made by entities that would not otherwise qualify for access (see below).

In both the agency-driven and peer-review tracks, allocations should be provided in credits with base rates derived from the cost of computational time or data storage. Some services, such as downloading data or models from a repository, would not require any credits.

The tracks should be structured with different criteria and processes for selection. Within the peer-review track, the basic principle would be that, as the size of the request grows larger, the bar for review increases. At the startup request level, the application would be a simple form that validates enrollment and eligibility, along with a description of the project. At the research request level, the application should be more extensive, including a proposal describing the work, underlying funding support, estimates of the computational resources needed, and so forth.

Possible outcomes include full acceptance, full rejection, cuts in the amount awarded, or re-directing the investigator to different resources within the NAIRR. For government-owned or -controlled resources made accessible to researchers through the NAIRR, the NAIRR resource allocation process should not bypass existing access approval processes but rather route NAIRR researchers into these existing processes.

If sufficient NAIRR resources are available, the Operating Entity may develop a direct-charge model for a subset of available resources. This “purchase” option can be useful both for granting access to users who would not otherwise be eligible for NAIRR access, as well as allowing those users who receive NAIRR access to grow their allocation beyond what can be freely provided. Revenues from cost recovery can be used to further expand the capacity of the NAIRR, providing access for additional users without sacrificing the availability of resources for the typical user base. A thoughtful and publicly-disseminated approach to establishing cost models can ensure that the NAIRR’s public funding stays consistent with the original goal of democratizing access.

The Operating Entity should establish an allocation system to award credits in alignment with available resources. Because AI workloads are extremely difficult to estimate in advance, NAIRR policies should permit the augmentation of resources through justified supplements, advances, or transfers from other accounts. The Operating Entity, with guidance from the Steering Committee and Program Management Office, should regularly review and adjust the division of resources across the agency-driven and peer-review tracks.

Privacy, Civil Rights, and Civil Liberties Protections

The NAIRR should serve as an exemplar for how transparent and responsible AI R&D can be performed with proper training and oversight at multiple levels. Processes to ensure that NAIRR operations, research, and governance are conducted in a transparent fashion with appropriate oversight should be integrated across all aspects of the design, implementation, administration, management, and use of the NAIRR. The NAIRR Operating Entity, with input from the advisory boards, must be proactive in addressing privacy, civil rights, and civil liberties issues. It must integrate appropriate technical controls, policies, and governance mechanisms from the beginning. One important initial step will be to include a diverse set of experts from relevant disciplines as part of NAIRR leadership and governance. The Steering Committee, Program Management Office, and Operating Entity must work together to ensure diversity among NAIRR decision-makers, and draw from the expectations for automated systems described in the Blueprint for an AI Bill of Rights¹⁸ as well as best practices defined in the AI Risk Management Framework (see Box 5). The Operating Entity leadership should hire staff with expertise in protecting privacy and mitigating ethical and societal issues, who would work with the advisory boards to design privacy, civil rights, and civil liberties considerations into the Operating Entity’s governance and review structures and activities.

Consideration for ethical issues should be foundational to the NAIRR and permeate its decision-making processes. One specific area for attention is the data to be incorporated into the NAIRR. The Operating Entity should develop publicly reviewable controls for datasets that the NAIRR hosts and a mechanism to ensure that datasets with legal, ethical, or discriminatory issues are quarantined and appropriately handled, drawing from the principles and expectations detailed in the Blueprint for an AI Bill of Rights. This should include support for system auditing and for maintenance of an archive of retired datasets to provide researchers with the ability to study data with different types of biases to better understand common data issues and potential harms, as well as the robustness of AI models when applied to such datasets.

The Operating Entity should establish, implement, and publicize acceptance criteria and recommended best practices for all resources joining the NAIRR to ensure that they are vetted from privacy, civil rights, civil liberties, and equity perspectives. These acceptance criteria should be more stringent for resources that are likely to be used in contexts that raise heightened concerns about privacy, civil rights, and civil liberties. It will be critical for the NAIRR to act quickly to provide such information, because much harm can result from delaying decision-making.

The impacts of any controls instituted should be evaluated and adjustments made as needed. The Ethics Advisory Board, in consultation with the User Committee, should play a central role in designing and implementing privacy, civil rights, and civil liberties requirements across all NAIRR systems, policies, and practices, and in ensuring dissemination of those requirements across the ecosystem. The uptake and use of the requirements should be incorporated into the NAIRR KPIs.

The Operating Entity should work with the Ethics Advisory Board to develop criteria and mechanisms for evaluating research and resource proposals from a privacy, civil rights, and civil liberties perspective; submit these criteria and mechanisms to the Program Management Office for review by the Steering Committee; and publish the criteria on the NAIRR website.

Box 5. Guiding Principles for NAIRR Policies

Multiple efforts are underway nationally and internationally to articulate responsible AI principles and operational strategies. The Blueprint for an AI Bill of Rights was released by the White House in October 2022, and includes a set of five principles and associated practices to help guide the design, use, and deployment of automated systems to protect the rights of the American public in the age of artificial intelligence.¹⁸ These five core protections are: safe and effective systems; algorithmic discrimination protections; data privacy, notice, and explanation; and human alternatives, consideration, and fallback. The Operating Entity should consider this framework when developing its policies and procedures.

NIST is developing an AI Risk Management Framework, which is anticipated to be released in early 2023. The framework is being developed through a consensus-driven, open, transparent, and collaborative process, and compliance will be voluntary.¹⁹ Overall, the framework is intended to give AI developers the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems. The Operating Entity should consider this framework when developing its policies and procedures.

Finally, ensuring awareness about rights, responsibilities, and best practices related to privacy, civil rights, and civil liberties is essential. All NAIRR users will be required to complete training, renewed annually, before being granted access to the NAIRR.

Scientific Integrity

The Operating Entity should also be responsible for addressing scientific integrity concerns. The Operating Entity should work with the User Committee to develop criteria and establish mechanisms for addressing researchers' and AI users' concerns associated with NAIRR-enabled research, submit them to the Program Management Office for review by the Steering Committee, and publish the criteria on the NAIRR website. These criteria and mechanisms should be informed by the Presidential Memorandum on Restoring Trust in Government Through Scientific Integrity and Evidence-Based Policymaking²⁰ and the guidance put forward in the *2023 Framework for Federal Scientific Integrity Policy and Practice* from the National Science and Technology Council's Scientific Integrity Framework Interagency Working Group.²¹ There should be mechanisms that allow early, easy, safe, and confidential reporting of perceived concerns. The Operating Entity staff should work closely with the Ethics Advisory Board to ensure that best practices are followed and that concerns are quickly addressed. KPIs should be established to ensure that this goal is satisfactorily met.

The Operating Entity should provide public information about research performed using NAIRR resources through regularly updated and publicly available project registries containing information such as (1) project names, descriptions, and anticipated value to the public; (2) project teams and affiliations; (3) data used; (4) research questions and methods; and (5) anticipated deliverables and associated delivery dates. The processes and policies established by the Operating Entity should reinforce the expectation that data, code, and publications resulting from federally funded research should be made publicly accessible to the extent possible. Users would be expected to comply with Federal agency public access policies updated in response to the memorandum issued by OSTP on August 25, 2022.²²

System Security and User Access Controls

The cybersecurity threat landscape is rapidly changing and evolving as new actors, attack methods, and vulnerabilities emerge. AI research, as an asset to economic growth and national security, is a high-value target. Cybersecurity risks extend beyond technical considerations to human behavior. Creating a culture of usable security and training is key to mitigating human mistakes that can lead to compromise. Just as convenience could conflict with security, fostering an open research environment has tradeoffs with providing secure access to high-value information and resources and protecting intellectual property.

The Operating Entity should implement system safeguards using government-applicable NIST security guidelines as well as the Five Safes framework: safe projects, safe people, safe settings, safe data, and safe outputs. The Five Safes framework structures protection across five dimensions: research projects and individuals working on projects are reviewed and approved; people using the resource must sign security agreements and complete training, and users' access is monitored; settings operationalize security needs and are managed through a central platform;

data is appropriately safeguarded against security, re-identification, and privacy risks; and exports are technically and contractually controlled, and evaluated and monitored to prevent unauthorized disclosure.

The Operating Entity should design the NAIRR to consist of multiple tiers, starting with at least two primary zones: an open science zone, NAIRR-Open, and a secure zone, NAIRR-Secure. Each zone will federate computational, network, and data resources operating in accordance with security and access-control policies that are uniform within the zone, but different between zones, reflecting the different priorities and needs of the users and resource operators. For example, ease of access and use may be of greater importance in the open science zone and appropriate for classroom settings, while data security may be of greater importance in the secure data zone and appropriate for sharing and analyzing Federal agency protected data.

The NAIRR-Open zone should adopt the best practices developed over two decades in the open science community, drawing from experiences and approaches used by ACCESS, the Open Science Grid, and the National Research Platform.²³ Access to open science resources should be managed using single sign-on authentication and a resource allocation mechanism managed by the Operating Entity.

The NAIRR-Secure zone should consist of one or more secure enclaves adhering to a common set of security controls,²⁴ and have the ability to support security requirements for sensitive information, such as those necessary to protect Controlled Unclassified Information and those arising from the Health Insurance Portability and Accessibility Act and other laws and regulations.²⁵ User-based access will be an important element in the NAIRR-Secure zone. The NAIRR-Secure zone should be administered by a specialized resource provider, subject to all of the oversight and reporting responsibilities of any NAIRR resource provider, but with the additional responsibility of security monitoring and controls compliance for its set of managed projects. To the extent that the data owners (e.g., Federal agencies, other non-governmental resources) require an Authorization to Operate, then it will be the responsibility of the NAIRR-Secure resource provider to obtain it.

Because the datasets to which the NAIRR provides access could include sensitive data on human beings or confidential government data, and because the security landscape is constantly changing, the Operating Entity will require staff with expertise in security, privacy, and usability, and will need to establish security controls and mechanisms that can keep up with the rapid pace of change and ensure the security and confidentiality of such data in accordance with Federal regulations. The value of access to sensitive data is also constantly changing, as evidenced by the recent experience with the COVID-19 pandemic; as a result, the Operating Entity will also require staff with expertise in measuring the value and use of data access, in accordance with the requirements of Title II of the Evidence Act. The Operating Entity must also comply with all Federal regulations for protected data, and adopt both value- and risk-based approaches for protecting sensitive data not otherwise covered by Federal regulations.

Open-Source Principles

The NAIRR Operating Entity and resource providers should adopt the principle of open source for products developed with Federal funds. Exceptions to open-source requirements should be provided for small businesses supported through SBIR or STTR programs that are given access to the NAIRR, and in cases where data are protected. The Operating Entity should leverage existing programs at Federal agencies that support translational activities such as having a professional software developer package software and tools developed as part of research projects for longer-term open-source availability. The NSF Cyberinfrastructure for Sustained Scientific Innovation (CSSI)²⁶ and Pathways to Enable Open-Source Ecosystems (POSE)²⁷ programs are two relevant examples of existing programs that focus on open-source development and support such translational activities.

More generally, research products should be made freely available through the NAIRR so long as they are reasonably mature and documented (i.e., production-level resources).

Environmental Sustainability

A system to source hardware in an environmentally sustainable way and measure and manage discarded hardware and other electronic waste (i.e., electronic devices that have reached their end of life) should be established for all resources made available through the NAIRR. Key elements of electronic waste management include maximizing the life cycle and usability of systems, as well as plans for electronic waste recycling, systems and equipment repurposing, and hardware reselling. Recycling electronic waste presents an opportunity for the recovery of critical minerals, in addition to reducing greenhouse gas emissions and limiting disposal. When reuse or recycling is not possible, disposal of electronic waste should involve accurately characterizing the waste and sending it to proper permitted disposal sites. For all discarded equipment, records should be kept tracking the disposal of potentially hazardous waste.

The Operating Entity, with the assistance of its Technology Advisory Board, should also work toward identifying computing technologies that are energy efficient and carbon neutral, and that have little or no negative effect on water quality, air quality, waste accumulation, soil contamination, or the U.S. carbon footprint. The Operating Entity could consider evaluating potential resource providers based on the energy efficiency and/or environmental sustainability of the design of the proposed resources. For example, resource providers could work with the Environmental Protection Agency's Energy Star for Data Centers program²⁸ to improve efficiency, reduce data center cooling energy, and optimize environmental performance.

The Operating Entity and resource providers should acquire, develop, and promote the use of tools to monitor and optimize applications for energy-efficient operation. This would require NAIRR resources to be instrumented with technologies that can identify utilization and energy use at the component level, as energy usage is specific to an application's execution. They should also identify application development tools and environments that can assist a programmer in the creation of highly energy-efficient applications and promote energy-efficient user behaviors. These tools should also help the operating system to allocate system capacity to each application with the goal of optimizing energy use.

The Operating Entity should also promote the importance of studying environmental issues through its support of relevant AI research areas. It should track and report on the percentage of time the NAIRR infrastructure is used for environmental research. Possible areas of study include environmental systems modeling and analysis, climate modeling, bio-systems modeling, watershed modeling and analysis, energy systems management, and waste management. Predictive maintenance and sensor systems learning are other relevant areas of AI research.

4. NAIRR Structure and Specifications for Resource Elements

The NAIRR Operating Entity should develop an integrated portal to provide the user base described in Chapter 2 with access to a federated mix of on-premise and commercial computational and data resources and services. Computational resources would include conventional servers, computing clusters, HPC, and cloud computing, and should also support access to edge computing resources and testbeds for AI R&D. The NAIRR Operating Entity should make open and protected data available via resource providers and partnerships. Data should be co-located with computational resources where possible. Data providers should facilitate user access to restricted statistical data through the Standard Application Process (SAP) established under the 2018 Foundations for Evidence-Based Policymaking Act, where appropriate and possible.²⁹ The NAIRR Operating Entity and resource providers should make software, training, and educational resources available to support a diverse set of users with varying levels of AI research experience and proficiency.

This chapter provides details of these key components, along with desired capabilities when the NAIRR begins initial operations. Given the fast pace of technological development, the Operating Entity should maintain the flexibility to adjust approaches to the elements detailed below, in consultation with the Steering Committee and Program Management Office.

Access Portal and User Interface

The Operating Entity is responsible for development of an NAIRR user portal that supports key user functionalities such as single sign-on, team allocations, data search and discovery, collaboration tools, resource discovery, job submission, consolidated accounting, spend alerts, information about data use, and cost-optimization of workflows. The portal will be one way to access NAIRR resources. Alternate access methods such as secure shell or scripting interfaces should also be made available for advanced users. The portal will allow users to select their AI applications, computational resources, and data sources from a curated catalog, and to launch and monitor jobs from a portal that provides a uniform, integrated view.

The portal should have built-in help functions and an integrated help desk ticketing system. The portal should maintain an up-to-date catalog of resource provider user documentation and training materials. Chat functions, meeting rooms, forums, and other functionality may be included to support collaboration and community building among students, researchers, resource providers, and other users. The portal should also enable data search and discovery and leverage automated technologies so that (1) metrics on data use can drive data acquisition and (2) diverse, community-driven data curation, linkage, and validation activities can be fostered. A user account would be required to manage computational allocations, monitor usage, submit jobs, and post to the community forum.

The Operating Entity should provide a public website through which some key elements are available without the need for a user account and sign-on. For example, linked catalogs of AI

education tools and testbeds, as well as an index of AI datasets with metadata, annotations of known problems and deprecation status, and community-contributed code, should be readily available.

The Operating Entity should assess the cost of building the user portal and public website in-house versus acquiring it commercially. To speed development, the Operating Entity could outsource the design, construction, and maintenance of the user portal to a commercial entity that has previously created successful user portals. All major aspects of the portal should be included in NAIRR initial operational capabilities.

Computational Resources

To lower barriers to entry into AI research, the Operating Entity and resource providers must make access to computational and data resources available to a variety of new users who otherwise would face financial, logistical, or capacity challenges engaging in the AI research ecosystem. Expanded access should be provided by leveraging existing resources in all sectors, augmenting the capacity of federally provided resources as appropriate, creating new research computing and data infrastructure to serve the AI R&D community, and providing financial support where needed. The NAIRR should also support the federation of user-supplied computing resources, testbeds, and sensors at the edge.

Capacity and Capability

When fully implemented, the NAIRR should address both the capacity (i.e., ability to support many users) and capability (i.e., ability to train the most resource-intensive AI models) needs of the AI research community. To meet existing capacity needs, the NAIRR should provide a mix of computational resources (i.e., on-premise and commercial cloud, dedicated, and shared resources) with a range of central processing unit (CPU) and graphics processing unit (GPU) options with multiple accelerators per node, high-speed networking, and sufficient memory capacity (i.e., at least one terabyte per node). The exact balance of computational resources will depend on the results of resource provider funding opportunities. Users should have the option of selecting which resources they would like to use through a range of mechanisms, including the user portal, direct command-line access, or optionally interactive “notebook”-like environments.

To meet users’ capability needs, the NAIRR system should include at least one large-scale machine-learning supercomputer capable of training 1 trillion-parameter models. This could be made available by leveraging an existing supercomputer or newly procured through a competitive bid process managed by the Operating Entity in consultation with the Steering Committee and relevant advisory boards.

NAIRR Software Resources

AI research has grown explosively through the development and dissemination of open source software (OSS) frameworks including TensorFlow, PyTorch, and their derivatives. Both these packages were developed by commercial entities and could have been kept proprietary. Instead, they were released as OSS projects, to the benefit of, and for further development by, the

AI research community. The success of these projects has inspired many other OSS projects and tools.³⁰

The Operating Entity, with advice from the Technology Advisory Board, should assess OSS packages most used by AI researchers and specify a standard software environment for the NAIRR federation.³¹ This software environment should be containerized as a lightweight virtual machine, and be supported across resource providers. Academic teams with their own on-premise servers would be encouraged to adopt the NAIRR federation standard. In addition, the Operating Entity should explore new AI workflow orchestration tools and templates for standard AI analysis tasks, such as `cnvrg.io`,³² which can meet the needs of industry researchers and might be suitable for adoption by the NAIRR federation.

Data and Datasets

The Operating Entity should provide a search and discovery service with metadata about the usage of all datasets. Such a service should be consistent with Section 202(c) of the Evidence Act. It should be designed to dovetail with the capabilities anticipated through development of a Federal data catalog, but extend beyond Federal data.

The Operating Entity should support data resource providers by either funding the creation of or providing continuing support to existing AI data repositories. In coordination with the Technology Advisory Board, the Operating Entity should publish interoperability guidelines for such data repositories, and encourage data repositories to compete to become NAIRR data resource providers. These guidelines should be informed by the Desired Characteristics of Data Repositories for Federally Funded Research developed by the National Science and Technology Council's Subcommittee on Open Science.³³ Having such repositories and datasets visible, searchable, and discoverable inside the NAIRR, as well as implementing mechanisms to track dataset use, are important to the success of the NAIRR.

NAIRR-Open and NAIRR-Secure zones should federate computational, network, and data resources operating in accordance with security and access-control policies that are uniform within the zone, but different between zones, reflecting the restrictions associated with the data in each zone. NAIRR-Secure should coordinate and collaborate with the program office designated by the Office of Management and Budget to oversee the SAP, and others as appropriate, in making available and specifying security and user access controls required for restricted (confidential) government and third-party data.²⁹ SAP is required by the Evidence Act to be the "front door" for accessing restricted data within the possession of Federal statistical agencies.

Dataset Acceptance Criteria and Metadata Standards

The Operating Entity should evaluate and characterize datasets into tiers, each with a different level of acceptance criteria. Examples include high, medium, and low levels of metadata; provenance; information about dataset usage, and the availability of persistent identifiers. The Operating Entity should ensure that each dataset is evaluated according to industry standards or best practices and that a determination is made on how each should be categorized. Where possible, such cataloging efforts should be aligned with efforts to develop a Federal data catalog.

The Operating Entity should not define dataset standards, as this area continues to evolve rapidly and would be best addressed by the community of users. However, the Operating Entity should provide a public-facing list of acceptable formats to ensure compatibility with resources and tools, encourage broader use, and leverage existing community-driven principles and standards such as those developed by the Research Data Alliance and NIST, among others. Regardless of category, substantive documentation should be provided with each directory or file containing data. The Operating Entity should also specify what it means for a dataset to be “analysis-ready” and categorize datasets accordingly. For example, an analysis-ready dataset should be in a structured format (e.g., a relational table or JSON³⁴ or Neo4j³⁵ formats) and should include details such as the semantics and provenance, information about the data-generation process, a data dictionary, related code, summary statistics for quality-assurance purposes, and information about how it has been used in previous analyses. Further, such a dataset should conform to standards in cases where datatypes are normally represented in a standard ontology (e.g., geographic information system [GIS] vector objects, gene ontology codes for molecules). Not all datasets need be in analysis-ready form. Some types of data or partial datasets are important or rare, and can be contributed with the goal that others can help transform them into analysis-ready data.

Role of the Operating Entity in Incentivizing and Curating Contributed Datasets and Other Resources

Since the quality of many AI models depends on high-quality training and test data, the Operating Entity should establish a data service that facilitates access to and additional use of existing curated datasets of value and interest to the NAIRR user community. Curation of AI data, models, tools, and workflows should be done by the user community in an AI data commons, facilitated by the NAIRR search and discovery platform. Such a community system, governed by terms of use as well as a review system, would facilitate data sharing and curation by members of the community. In the context of a commons model, researchers who contribute to the common good through data curation and code sharing, and whose contributions are recognized and valued by relevant communities, could be incentivized through high-profile NAIRR recognition and/or preferential access to NAIRR resources.

The NAIRR Operating Entity should test, on a trial basis, a service for searching for, discovering, and curating valuable external data as well as data generated with NAIRR resources. One option would be to contract with one or more commercial AI marketplaces to meet its users'

data curation needs. The “AI marketplace” is a powerful concept that has emerged in the commercial sector; it refers to the social and technical infrastructure through which the user community contributes, documents, and shares data, codes, and models. Contributions are validated and valued by the community, and community standards are enforced by the company managing the marketplace. Another option is for the NAIRR to develop its own “AI data commons” with attributes similar to a commercial marketplace. Such an option is likely to be preferable for the federally funded NAIRR. However, since both commons and marketplace options have merit, the Operating Entity should have flexibility regarding development of data curation services, and the services should be implemented on a trial basis and evaluated for efficacy by the Operating Entity in the first five years of NAIRR operation.

Substantial Operating Entity resources should be dedicated to technical support staff who can support community-driven curation efforts. Data users, contributors, and curators will require support to understand and meet the technical standards of NAIRR data repositories. Further, training and additional support will be critical to the integrity and quality of NAIRR datasets, and to protect privacy, civil rights, and civil liberties.

The NAIRR and Existing Federal Government Data

Federal agencies hold data that could fuel foundational, use-inspired, and translational AI research in domains such as transportation, healthcare, and natural hazards research. Sources of Federal agency data include statistical data, administrative data, and data from federally funded intramural and extramural research. While some of these datasets are already accessible to the public, many others are not.

Since Federal datasets could be highly valuable to AI research and advance national goals, there are three other Federal Government data efforts with which the NAIRR could engage. One is data.gov, which is a website that points to other resources containing information and data generated by agency or agency-funded projects. Most of the retrievable data on data.gov are in web or text form, which might be of interest to some NAIRR researchers. However, scientific numerical datasets are deeply buried in data.gov and not easily accessible. The Operating Entity and Program Management Office could work with data.gov to encourage additional contributions conforming to NAIRR data acceptance criteria, which should include measures of data use. Another is the SAP, through which researchers will be able to discover and apply for access to restricted data acquired by Federal statistical agencies through a single application process and portal.³⁶ Finally, the National Secure Data Service (NSDS) demonstration project, established by the CHIPS and Science Act of 2022, has the potential to complement the SAP and existing statistical agency efforts with additional capability for data acquisition, linkage, and protection (see Box 6 for more details).

The Steering Committee should facilitate the establishment of a NAIRR-Federal Interagency Council on Statistical Policy (ICSP) working group. This working group should collaborate to assess options for establishing a secure node for the purpose of enabling large-scale AI analysis of government data for statistical purposes. Where such resources are not intended to be made accessible via the SAP or the NSDS demonstration project, the working group should define the Confidential Information Protection and Statistical Efficiency Act-compliant data access protocols and controls. This NAIRR-ICSP collaboration should facilitate the provisioning of timely access for appropriate (i.e., approved) projects to restricted (i.e., confidential) government and third-party data.

Box 6. The National Secure Data Service Demonstration Project

The CHIPS and Science Act of 2022 includes a provision that requires NSF to create a demonstration for the National Secure Data Service (NSDS). The intent of this demonstration is “to develop, refine, and test models to inform the full implementation of the Commission on Evidence-Based Policymaking recommendation for a government-wide data linkage and access infrastructure for statistical activities conducted for statistical purposes.”

The NAIRR should also encourage and support additional contributions of State and local datasets conforming to NAIRR data acceptance criteria, and subject to the legal requirements of the State and local government agencies, either by working with data.gov³⁷ or the eventual NSDS.

In terms of existing high-quality data repositories managed by agencies such as the National Institutes of Health (NIH) and NASA, the Operating Entity will need to determine whether to reproduce large datasets that are already available from these other sources or find other means of coordinating access for NAIRR researchers. This coordination could benefit from regular convening of leadership from various Federal data efforts to identify ways to improve coordination and avoid inefficiency or redundancy.

Legal Compliance

The Operating Entity should ensure that data access through the NAIRR is in compliance with applicable Federal laws. Consider, for example, data use agreements (DUAs), which are contractual documents established between provider and recipient institutions and used for the transfer of nonpublic or restricted data. A DUA in the case of the NAIRR would benefit from being structured around the Five Safes framework to ensure safe use. Generally, a DUA will define publication responsibilities, disposition of intellectual property arising out of the use of the data, ownership of derived datasets, and expectations for disposal of the data. The use of a DUA is good practice because it establishes a clear understanding of the expectations and responsibilities of both parties.

It is anticipated that an SAP Governing Board²⁹ will be the primary Federal entity with responsibility for overseeing the process by which secure access to protected Federal statistical data is approved for both government and external users, taking into account aspects of privacy, civil rights, and civil liberties. Rather than create a duplicative infrastructure, the Operating Entity

should coordinate closely with the SAP Governing Board to ensure that NAIRR users are aware of and have appropriate access to Federal statistical data provided through the SAP.

Co-Location of Resources

AI training datasets can be many terabytes in size. With current technology, moving this volume of data over the commercial internet would take many hours at typical network speeds. Effective computing within a research cyberinfrastructure that handles high-volume data will likely require the co-location of data with the hardware on which it will be processed. The Operating Entity should facilitate the co-location of data and computational resources in two ways: (1) invest in the build-out of a NAIRR AI data commons infrastructure at the HPC centers coupled with an expansion of computational capacity and (2) negotiate contracts with the public clouds with educational discounts that provide access to the most popular computational and storage solutions for AI researchers. The Operating Entity should also provide access to existing AI-relevant resources that co-locate computation and data.

The Operating Entity should additionally create and curate a searchable and discoverable catalog of existing and available governmental and non-governmental datasets, including providing information about their usage, that may be distributed across the United States. These datasets, particularly the confidential data, need not be co-located with the computational resources provided by the NAIRR, although some datasets could be copied to co-located storage to facilitate better performance. Datasets created using the NAIRR infrastructure should be stored at co-located NAIRR storage facilities. Thus, there should be a mix of distributed and co-located datasets as part of the NAIRR infrastructure with multiple mechanisms to support efficient use of those datasets, including a partnership with the SAP Governing Board and eventual NSDS.

Educational Tools and Services

To lower the barriers to participation in the AI ecosystem and increase the diversity of AI researchers, the NAIRR must be broadly accessible to a range of users and include educational and technical information. The NAIRR access portal should provide catalogs and search and discovery tools to facilitate access to educational and training materials for a range of experience levels.

The NAIRR should provide a platform that can be used for educational and community-building activities. This platform can provide facilitation functions for educational efforts, but the Operating Entity should not be responsible for developing general or discipline-specific educational content, because general education on AI and computational expertise is not the primary mission of the NAIRR.

Technical training and support materials related to the use of the NAIRR are within scope, and the Operating Entity and resource providers should share the responsibility for training and support in the use of NAIRR resources.

Tiered Technical Training and Support

The Operating Entity should provide technical training materials for users at different skill levels (e.g., beginner, intermediate, and advanced). Training options should span a range of formats, including web pages, tutorials, webinars, online training, and customized remote workshops. Training should include use of the portal itself, in addition to training and other information on the particular resources available via the NAIRR portal, as well as NAIRR policies and procedures.

Curation of Training Materials

To support the needs of a diverse set of users, the Operating Entity should build a consolidated, searchable catalog of training materials generated by NAIRR resource providers so that everything is listed in one place. Resource providers should provide context-based training resources as well as just-in-time training. The Operating Entity should also facilitate identification and curation of additional AI- and resource-related training materials by the user community. The system should be instrumented to track highly used pages and tutorials to help resource providers better understand how users are getting the information they need and to refine how the content is delivered (e.g., static documentation versus interactive tutorials).

The level of training required should be commensurate with the nature of NAIRR usage. For example, short-term, non-sensitive use of the NAIRR, such as a short classroom exercise, may warrant less rigorous requirements. Because the user base for the NAIRR is intended to be broad and diverse, training should be tailored for various audiences. Tiered user training documentation (e.g., beginner, intermediate, and advanced) and interactive tutorials should be created and kept current by resource providers.

Platform for Educational Activities

The NAIRR should provide user access to educational infrastructure made available by educational resource providers. An example of this concept can be found in CloudBank, which provides users with access to the Berkeley Data Stack,³⁸ a collection of tools and resources that support data science research and education at the University of California, Berkeley. The Berkeley Data Stack provides each student with an interactive learning environment via a Jupyter notebook interface to Jupyter Books, integrating notebooks and computational content with textbooks developed by the instructor.

Technical Integration

Software for Integration

Software will be needed to federate the diverse resources incorporated into the NAIRR. Examples include grid toolkit software, an information-publishing framework, resource-description repository, accounting and account-management software, a common user environment, a single sign-on hub, and file transport services. As an example that the NAIRR could build from, many of these solutions are being used in the NSF ACCESS program (i.e., the follow-on to XSEDE, which began in September 2022).³⁹ The NAIRR should leverage such

developed software approaches, and the NAIRR Operating Entity (with advice from the Technology Advisory Board) should evaluate existing integration software stacks such as that used in ACCESS for possible adoption.

The NAIRR infrastructure should support distributed workflow orchestration software.⁴⁰ The NAIRR user portal will need to be fully integrated with these software functions as part of NAIRR's full operational capabilities.

Integrating Data Resources

One approach that will facilitate NAIRR technical integration is incorporating Federal data resources stored in commercial clouds. Several Federal agencies have placed large datasets of potential interest to external researchers in the commercial clouds, taking advantage of the public data hosting programs. A June 2022 National Science and Technology Council report entitled *Lessons Learned from Federal Use of Cloud Computing to Support AI R&D*⁴¹ notes that “use of the cloud has simplified computational access to data owned and maintained by Federal agencies, facilitating efficient use of and collaborative work with big data. For example, over 36 petabytes of public and controlled access genomic sequencing data hosted by the NIH's National Library of Medicine are now available on two commercial cloud computing platforms,⁴² and 10 petabytes of public weather and environmental data are now accessible through the National Oceanic and Atmospheric Administration (NOAA) Open Data Dissemination Program across three commercial cloud computing platforms.⁴¹ NASA has taken similar steps, storing newly collected Earth Science data in the cloud to make it easier for the public to access and reduce the requirement of downloading data to perform analytics.”⁴³ The Operating Entity should leverage and replicate this approach to enable effective use of large-scale data in the cloud.

Testbeds

AI testbeds are simulated, live, or blended environments that support research, prototyping, development, and testing of AI applications. Increasing access to testbeds via the NAIRR will provide researchers without institutional testbeds the opportunity to explore new approaches for solving important problems. Testbeds can be broadly defined as serving the purpose of either comparison or validation. Comparison testbeds allow researchers to measure the effectiveness of new engineering, math, or algorithmic developments. These testbeds can take the form of test frameworks and competitions, simulated environments, or living laboratories and are useful for foundational, use-inspired, and translational AI R&D. Validation testbeds allow developers to decide whether it is acceptable to move up the maturity cycle of an end-to-end system to a more advanced phase of development, and are useful for translational research. Note, however, that validation testbeds supported through the NAIRR are intended for early-stage and translational research, rather than for the purpose of validating commercial products.

The Operating Entity should facilitate connections to AI testbeds. It is likely that each AI testbed will have unique requirements for connection and/or integration. The Operating Entity, with consultation from the Science Advisory Board and Technology Advisory Board, should determine which testbeds should be made accessible via the NAIRR as part of initial operational

capabilities, including through consideration of which interfaces, protocols, and controls are necessary to facilitate access to each.

With an AI data commons model, testbeds can be reviewed and made available, maintained by their creators with the incentive of exchange with other assets in the marketplace. The Operating Entity should work with the Networking and Information Technology R&D (NITRD) Program, which catalogs Federal AI testbeds, to expand the inventory beyond federally funded resources. NITRD may wish to transfer this responsibility to the Operating Entity.

5. Phased Buildout of NAIRR Organization and Resources

The NAIRR cyberinfrastructure should be established in a phased manner with a gradual ramp-up of resources over time. Phasing can help ease the process of integration across the federated NAIRR system, provide opportunities for users to transition as older resources age out and new resources come online, provide value to users more quickly, and allow the NAIRR Operating Entity to receive user feedback expeditiously.⁴⁴

This approach is also intended to avoid challenges associated with acquiring AI-relevant cyberinfrastructure, which develops at a rapid pace and can quickly become outdated. Agencies that have already invested in AI should be part of a collaborative process for identifying the computational, data, and training needs. New agencies that are just beginning to invest in AI can work with other agencies to identify gaps and capabilities that would be useful for those agencies' missions.



NAIRR implementation has been divided into four phases, as indicated in the graphic above. The timelines in this report assume that work will begin immediately after the publication of this final report, but they may also be adapted as appropriate. To start, the federated NAIRR system should be built out from the baseline of existing computational and data resources, augmenting their capacity and capability while making them discoverable and accessible through the NAIRR user portal. This should be accomplished in parallel to investments in new computational and data resources to serve and grow the capacity of the AI research community. A NAIRR Pilot Option could run in parallel to this buildout, as described in Box 7.


NAIRR should achieve initial operational capability—availability of the core user portal and a basic complement of computational and data resources for users—no later than 21 months from the U.S. Government launch of the program. Steady-state operations, during which the

Box 7. NAIRR Pilot Option

The implementation plan presented in this report targets an initial operation of the NAIRR in late year 1. To expedite the availability of AI research resources to the AI R&D communities as early as year 0, the NAIRR Task Force proposes that the NAIRR Program Management Office provide pilot-scale access to existing computational resources, software, datasets, services, and user portals across the current national cyberinfrastructure ecosystem, by providing supplemental funds for this additional use by the beginning of year 1 and issuing broad calls to the AI R&D community to apply for this access. Setting up such a pilot would require rapid establishment of interim management and governance mechanisms. The pilot would operate until the NAIRR is fully operational in year 2, at which point it would ramp down; the Program Management Office can incorporate its learning from this experience into its implementation of the NAIRR.

cyberinfrastructure system has met target capacity and capabilities for all components, should be established by the fourth year, with the understanding that the system should evolve and grow on an ongoing basis. Periodic evaluation and horizon scanning should inform changes to system operations, governance, and technology components to keep the federated infrastructure current and optimize utility.


Phase 1: Program Launch and Operating Entity Selection



The first steps to launching the NAIRR are the responsibility of the Federal Government. Congress should authorize and appropriate funds to establish the NAIRR as soon as possible. The NAIIO within OSTP, together with the agency that serves as the administrative home for the NAIRR Program Management Office, should coordinate the formation of the Steering Committee, and the agency that serves as the administrative home should stand up and staff the Program Management Office. The Program Management Office and the Steering Committee should write and release the funding opportunity for the Operating Entity within the first six months and establish the criteria and process for selecting the awardee. The Steering Committee should work toward developing necessary coordination processes for the selection and funding of NAIRR resource providers.

During months 6–12, proposals for management of the Operating Entity should be received, reviewed, and decided on by the Program Management Office, under the oversight of the Steering Committee, using the defined selection process and criteria. By the end of this period, the contract for the Operating Entity should be made, and the awardee should begin work.

Phase 2: Operating Entity Startup



Internal Planning and Operations

The Operating Entity startup phase begins when the contract has been established. As soon as possible, the Operating Entity should hire staff; establish strategies, policies, and procedures; charter and stand up the User Committee and advisory boards, establishing the Ethics and Technology Advisory Boards as soon as possible, and the Science Advisory Board within six months of the award; and conduct information-gathering and assessment to inform the design of the NAIRR user portal, interface, security and access controls, and support services. The Operating Entity should build in technical and policy tools to support privacy, civil rights, and civil liberties

considerations and NAIRR evaluation and assessment planning into its policies and procedures; these considerations must begin as soon as possible. In its first six months, the NAIRR Operating Entity should initiate biannual (or more frequent as needed) meetings of its boards and committees, develop governance policies and legal frameworks for constituent participation, and develop business processes and policies.

Within six months of its award, the Operating Entity should have developed and published necessary operational plans and policies, with input from the Program Management Office, Steering and User Committees, advisory boards, and other constituents—including members of the public and public interest groups. These include operational plans for the following:

- (1) Addressing privacy, civil rights, and civil liberties issues.
- (2) Creating NAIRR scientific integrity policies, user policies, data use agreements, and other legal requirements.
- (3) Developing specific user access controls and security architectures for both NAIRR-Open and NAIRR-Secure.
- (4) Supporting the process for selection of NAIRR resource providers.
- (5) Incentivizing participation and resource contribution, including through establishment of an AI data commons.
- (6) Managing resource allocations and user onboarding, including procedures for soliciting, reviewing, and managing those research proposals for which it directly administers resource allocations, and coordinating with agencies on allocations reserved for agency-funded researchers.
- (7) Providing transparent communication of information about how to access resources via the NAIRR—along with catalogs of AI resources such as datasets, software, educational tools, and testbeds—through a public-facing website.
- (8) Gathering and providing information to the independent, external evaluator, to ensure that NAIRR performance assessment can be planned early and infrastructure elements can be designed and adapted to facilitate collection of key data for assessment of KPIs across all NAIRR operational stages.

These plans should be reviewed periodically over the life cycle of the NAIRR and adapted as needed for different phases of operation and to best achieve the NAIRR’s KPIs. Work should be focused on meeting strategic objectives and goals as the research community needs evolve over time.

Establishment of Initial NAIRR Resource Components

In its startup phase, the Operating Entity should federate the first resource providers, establish an appropriate portal and user interface for accessing these resources, and identify its external evaluator in coordination with the NAIRR Program Management Office. As part of these efforts, the Operating Entity, Program Management Office, and Steering Committee should develop coordinated, multi-agency funding opportunities for resource providers as soon as possible, ideally

within six months of the initial Operating Entity award. These opportunities should be funded by the Steering Committee agencies and administered by the Program Management Office. The funding opportunities should also (1) call for the inclusion of existing resources that could be incorporated into the NAIRR without the need for additional funds and (2) fund the expansion of AI-capable computational and data resources at a subset of competitively selected existing advanced cyberinfrastructure sites. In addition, the Operating Entity should negotiate one or more public cloud contracts at discounted rates to provide researchers with access to the latest technologies and cloud-resident datasets with minimal startup overhead.

Winners of the funding opportunities should be chosen based on the scientific and technical merit of the proposals, cost effectiveness, and the suitability of the proposed systems for advancing and democratizing AI R&D. The first round of funding opportunities should allow additional time (not repeated in future opportunities) to bring the resources to a production state, as the technical integration process might still be under development for the first cohort. Subsequent opportunities should be used to fund the procurement and operation of new AI-tailored resources, both experimental and production, cloud and on-premise, and shared and dedicated, at new or existing sites.

Staffing for user support should be included in the proposal of any resource provider. Resource providers should be expected to provide competent technical support for users of the resources they provide, although the Operating Entity staff should provide help-desk functions. User-training materials should be developed and made available before the launch of the infrastructure. A separate resource provider for curation of education and training materials and catalogs of testbeds and datasets (with metadata including history and deprecation status) could also be funded if the Operating Entity does not manage this in-house.

The overhead cost for an open-data system is dramatically lower than that of a system that holds sensitive data; the legal and user agreement requirements are less stringent for open data as well. Both types of data will be necessary for a successful NAIRR, and providers will need to be identified and funded if the Operating Entity does not develop this infrastructure in-house. Open data can probably be made available prior to sensitive data, even if the resource providers begin work simultaneously. The initial set of opportunities for data-resource providers should include both providers of open data and providers of secure access to sensitive data. However, it is reasonable to expect that the initial roll-out to users will support only open data, because the legal and regulatory issues associated with sensitive data likely require more time to address.

The Operating Entity should determine its approach to the design, construction, and maintenance of an integrated user portal and interface to all resources that are part of NAIRR-Open, establishing preliminary capabilities during the startup phase. The Operating Entity should also immediately invest in building an evaluation data infrastructure sufficient to establish benchmarks and track progress over time. The evaluation data should include internal data about awarded and declined research proposals, as well as resource allocation information from all participating agencies. Information about the publications and patents resulting from research and researchers leveraging the NAIRR should be captured using automated methods. Administrative data from Federal, State, and local government data sources, as well as the private sector, should

be used to capture economic impact, leveraging the National Secure Data Service where applicable. (See Appendix H for illustrative examples of potential KPIs or evaluation metrics.) An example of a successful approach is the Institute for Research on Innovation and Science (IRIS) at the University of Michigan.⁴⁵

All of the Operating Entity’s startup activities should leverage the support of the User Committee and advisory boards, to the extent possible, to gather information and assess R&D community needs. The Operating Entity should consult frequently with the NAIRR Program Management Office and Steering Committee throughout its startup activities.

Phase 3: NAIRR Initial Operational Capabilities



The goal of the initial operational phase is to establish policies, processes, and technical resources that can be accessed in the near term by AI researchers and developers and that will support further buildout and maturation of the NAIRR. Initial NAIRR operational capabilities should be made available to researchers within nine months of the Operating Entity award. These capabilities should consist of (1) a portal and associated user-support resources, including indexes of resources and training materials; (2) a mix of operational on-premise and cloud resource providers, preferably with access to at least one ML supercomputer capable of training one trillion-parameter models; (3) a workable allocation and identity-management system; and (4) a workable data-publication system that allows datasets to be added to a catalog with a digital object identifier. These elements are sufficient at launch, although there are more that should be added soon after (e.g., common software stack, automated monitoring, AI data commons). The NAIRR-Open portal and at least some data sources should also be available.

The NAIRR’s initial operational capability should include a minimum complement of resources for users in the near term. The NAIRR-Secure portal and enclave, sensitive datasets, and new experimental and production AI-tailored computational resources may require additional time to mature and enter use. These resources should continue to develop during initial operations, with the goal of bringing all first-cohort resource providers into operational use by the end of the initial operating phase.

Initial Computational Resources

To facilitate a federation of existing on-premise and commercial cloud resources, established Federal agency programs could be leveraged. For example, the NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) Initiative program provides access to Amazon Web Services, Microsoft Azure, and Google Cloud resources.⁴⁶

The NSF pilot CloudBank⁴⁷ provides a portal with four commercial cloud resources (i.e., Amazon Web Services, Google Cloud, Microsoft Azure, IBM Cloud). The NSF Partnership to Advance Throughput Computing⁴⁸ and National Research Platform²³ programs provide access to federated national infrastructure including commercial cloud services.

While it might not be feasible for the initial resource mix to provide a full complement of architectures, it should include at least one “experimental” resource with something other than common CPU/GPU hardware (e.g., embedded or Internet of Things infrastructure, or new silicon for AI). For example, NSF and DOE support several federated distributed computing infrastructures that facilitate emerging AI experiments in a cyber-physical environment. When crafting the multi-agency funding opportunities for new resource providers, the Steering Committee, Project Management Office, and Operating Entity should consider how to complement existing federally supported resources included as part of the NAIRR initial operational capabilities.

Based on a Task Force analysis, the computing capacity goal for the NAIRR is 48–60 million hours on quad-GPU nodes in its initial operational capability. This level of capacity would allow 50,000 researchers (including students) to have access to 1,000 hours per user. Alternatively, 25,000 researchers would have access to 1,000 hours per user and up to 40 teams per year could solve a problem at the scale of OpenAI’s GPT-3 benchmark,⁴⁹ one of the largest (and most expensive to train) deep-learning models to date. By the time NAIRR reaches its full operating capability near the end of year 2, three times this capacity (140–180 million hours on quad-GPU nodes) should be available. This capacity corresponds to a NAIRR steady state supporting 150,000 researchers to have access to 1,000 hours per user of computing time, or, alternatively, 75,000 users and up to 120 teams per year could research GPT-3 benchmark-level problems.

Initial Data Resources

Data resources made available during the NAIRR’s initial operations should leverage existing Federal and commercial data repositories. Particular attention should be paid to existing large-scale data infrastructures that have co-located data and computational resources, such as the NIH All of Us⁵⁰ and National COVID Cohort Collaborative⁵¹ programs, to develop approaches for linking them into the NAIRR in alignment with the interoperability guidelines developed by the Operating Entity.


An initial instantiation of the NAIRR AI data commons should be in place during the NAIRR’s initial operations, with infrastructure and staff support for data hosting, data search and discovery, sharing, and community curation activities. Work to develop the tiered-access infrastructure to enable the provisioning of approved access to sensitive data should also be underway, in coordination with the establishment of the SAP and NSDS.

Initial Research Using the NAIRR

Once the NAIRR is available to support research, the Operating Entity should initiate processes to onboard users—both investigators who propose projects that would receive NAIRR resource allocations from the Operating Entity directly and those who are funded by and receive NAIRR credits or approvals from participating Federal agencies. To do so, the Operating Entity

must implement its policies and mechanisms for allocating computational credits and providing user access and training. During the initial operations phase, the Operating Entity should monitor system performance and resource utilization to learn lessons that can inform full operations. The advisory boards should review initial operations and provide advice to the Operating Entity regarding how the NAIRR can improve its performance during full operations.

Phase 4: NAIRR Ongoing (Steady-State) Operations



Evolution of NAIRR Resources

The NAIRR should evolve through periodic funding opportunities, developed in response to user uptake and demand. The first round should result in the selection of approximately one-third of the expected steady-state capacity of the NAIRR. Subsequent funding opportunities should be announced every other year at the same funding level. The Operating Entity should continue to solicit production and experimental resources and strive for architectural and resource diversity. Capability should be added to support emerging areas of interest and need by the research community and industry.

By three years after the initial operational capabilities are available, the set of resources needed to achieve full NAIRR capacity should have been funded by participating agencies, and their federation managed by the Operating Entity. Thereafter, approximately one-third of the resources may be replaced or updated every two years, while two-thirds of the resources should remain in production operation, providing both continuity and the opportunity to incorporate innovations. A minimum of 18 resource providers should be part of the NAIRR in the steady state, across a balance of types and architectures. Beginning in year 6, resource providers should be allowed to recompete.

Full NAIRR capacity should include a vibrant AI data commons as well as access to sensitive data through a secure, tiered-access system for vetted users and approved research projects. The Operating Entity, in partnership with the Steering Committee, should work across the Federal Government to make existing data repositories searchable, discoverable, and accessible via the NAIRR.

The Operating Entity should continue to take input from the research community via its User Committee and determine what capabilities should be added to the NAIRR infrastructure over time. These additions should be vetted through the Science Advisory Board and Technology Advisory Board, and inform the Operating Entity and the Program Management Office’s development of new funding opportunities and decommissioning of older components. All new capabilities should be added to the catalog of available infrastructure elements and made accessible via the user portal. The Technology Advisory Board should also periodically survey the evolving AI tool landscape and provide advice on additions or deletions from the NAIRR standard virtual

machine. The goal should be to add capabilities to continually serve the AI R&D needs of the NAIRR user community over time.

Partnership Engagement Operations

Partnerships will be important for providing the resources and expertise needed to maintain a cutting-edge NAIRR. The Operating Entity should establish public- and private-sector partnership mechanisms to extend both the NAIRR’s scope and its user base. Relevant partnership mechanisms would likely vary by sector and entity type. Since the NAIRR is envisioned as a large federation for a defined user community, partnerships could hinge on the balance of benefits for the NAIRR user community and the Nation writ large, and for the partner resource provider. In some cases, this would mean contribution of resources—either co-funding or in-kind support—in exchange for access to resource allocations to the partner or other benefits such as workforce recruitment pipelines, opportunities for collaboration, or to learn from NAIRR user research that leverages the resources provided. In the case of a public university, this might mean adding resources from a campus compute cluster or campus data collections to the NAIRR federation in exchange for additional time for that university’s researchers.

Private-sector partnerships could work similarly. Private entities can compete to become resource providers within the NAIRR, in which case they would make resources available in exchange for funding. But other models could also be defined, where companies could make in-kind contributions (e.g., tools, data, models, computational resources) in exchange for access to NAIRR resources.

User Outreach, Engagement, and Support Operations

The Task Force envisions a NAIRR where resource providers deliver support for the resources that they provision. Central operations functions, including support for the central portal and technical interoperation of the resource providers—as well as significant efforts for broadening participation, outreach, education, and training—should be the responsibility of the Operating Entity.

Outreach and International Collaboration

The Operating Entity should establish a small team to represent the NAIRR organization at international conferences where constituents gather—AI conferences as well as domain-specific areas in the life sciences, physical sciences, and social and behavioral sciences, etc.⁵² The team should document successes in science stories, ensure that NAIRR opportunities are disseminated broadly through domestic and international networks, and coordinate presentations and outreach in key forums.

Once the NAIRR has reached full operations, the Steering Committee and Operating Entity should explore ways to leverage the NAIRR to advance AI research through international cooperation with similar resource infrastructure efforts around the world. In doing so, the Operating Entity should follow the guidelines for international collaboration set by OSTP and U.S. Government research funding agencies, and comply with relevant export controls. The Operating Entity must also avoid activities or assigning access to its infrastructure to any embargoed or

sanctioned countries, institutions, organizations, or persons. Otherwise, the Operating Entity should work to establish collaboration and the sharing of information between U.S. and non-U.S. research entities. As it matures, NAIRR should leverage existing international forums such as the International Science Council’s Committee on Data (CODATA⁵³) and the Global Partnership on AI to support ongoing international collaborations and foster new opportunities.

NAIRR Budget

Based on a Task Force analysis of the estimated number of users and recent historical information regarding the cost of high-performance computing capacity, NAIRR costs are estimated at \$2.6 billion over a six-year period (see Table 1).⁵⁴ In estimating the budget for the NAIRR, the Task Force (1) focused initially on only the advanced computing resources that would be provided by (or through) the NAIRR, based on costs for existing advanced computing resources, and then (2) supplemented that estimate with estimates for other requisite NAIRR capabilities such as data, software workflows, and education and training. The Task Force assumed that all federally funded AI researchers throughout the United States from the targeted user communities would use the NAIRR to some extent. The Task Force further assumed that the average computing used by a NAIRR user would be comparable to that of a typical researcher using advanced computing resources. For additional context on the cost of training large ML models, see Box 8.

Table 1. NAIRR Six-Year Budget Summary

Year	Resource Providers	Operating Entity	Evaluation	Total
1	\$375M	\$70M	\$5M	\$450M
2	\$375M	\$60M	\$5M	\$440M
3	\$375M	\$60M	\$5M	\$440M
4	\$375M	\$60M	\$5M	\$440M
5	\$375M	\$60M	\$5M	\$440M
6	\$375M	\$60M	\$5M	\$440M
6-year total	\$2.25B	\$370M	\$30M	~ \$2.6B

Specifically, using one agency’s current advanced computing investments during the period FY 2016–FY 2021 as a proxy and considering known oversubscription of about 125 percent, the Task Force identified that an investment of over \$1 billion would have been necessary during this period. These investments would have provided advanced computing resources to a community of about 19,000 users spanning about 2,300 active projects totaling about \$6 billion in Federal R&D investment. Put another way, the average advanced computing investment needed per 1,000 users is about \$53 million, and the average advanced computing investment needed per \$1 billion of Federal R&D funding is about \$169 million.

Box 8. Training Large AI Models

Many recent breakthroughs in AI capabilities have been achieved through the creation of large, computationally-intensive deep learning models. In the pursuit of more generalizable capabilities, such models have been growing in size: OpenAI's GPT-3 in 2020 broke barriers at 175 billion parameters. Google followed suit in 2021 with a 1.6 trillion-parameter model, and the Beijing Academy of Artificial Intelligence with a 1.75 trillion-parameter model soon after. Published cost estimates ballpark that training a 110 million-parameter language model costs about \$50,000, a 340 million-parameter model costs about \$200,000, and a 1.5 billion-parameter model costs about \$1.6 million.⁵⁵ Overall, the cost depends on multiple factors, including size of the training dataset, model architecture, and the number of training runs.

In arriving at the final budget estimate, the Task Force took the above investments and estimates into account, along with an assumption that the scale of investment and size of the AI community will continue to grow rapidly in the years ahead. The bulk of the estimated budget of \$2.6 billion (i.e., \$2.25 billion) funds the NAIRR resource providers. Resource providers should be brought online every two years with a six-year lifetime, requiring a new \$750 million investment to be made every two years to ensure that NAIRR resources remain at the state of the art. The Operating Entity budget is estimated at approximately \$60 million

per year to support the coordination and management of NAIRR activities (see Table 2). An additional \$5 million per year is needed to support the Operating Entity's external evaluation process. The budget for the Operating Entity is based upon historical experience that the annual cost of operations for complex cyberinfrastructure is approximately 20 percent of the cost of the cyberinfrastructure resources themselves. Funding for the Operating Entity and external evaluation should be appropriated by Congress to the administrative home of the NAIRR, with suitable language to permit funds to be used to initiate and staff the Program Management Office. Funding for the NAIRR resource providers should be appropriated by Congress to the agencies that will fund them.

Resource providers should receive awards that allow them to provide services for up to six years to the NAIRR user community. Resource providers can fall into several categories, and in some cases the operation and acquisition costs for the resource may be blended. For example, a resource provider may have an initial cost in the acquisition phase for the hardware, followed by an annual cost in an operational phase to cover support personnel, maintenance, power, and so forth. A resource provider whose resource is providing training to the NAIRR may have almost no acquisition-phase costs and substantial operational costs. As a result, any funding opportunity for resource providers should include a mix of acquisition and operations funds for the resources themselves. Based on the experience of other federally funded computing operations, annual operations should not exceed 20 percent of the acquisition cost. Following this model, a six-year award should budget roughly 45 percent of the total for acquisition and 55 percent for operations. Resource provider awards should be capped at \$200 million, corresponding to a \$90 million acquisition with \$110 million for operations. To ensure a diversity of providers, the largest awards should be reserved for large computing investments, with smaller caps defined for data and service awards. A minimum of six awards should be made per cohort.

Table 2. Operating Entity Costs

Cost Category	Base Cost	Year One Startup Cost
Central Portal and Resource Integration	\$10M/year	\$15M
Training and User Support	\$15M/year	\$15M
Data Integration and Curation	\$5M/year	\$10M
Internal R&D and Technology Development Efforts	\$15M/year	\$15M
Other Operating Entity Allocations (e.g., advisory boards, governance activities)	\$15M/year	\$15M
Total	\$60M/year	\$70M

NAIRR Evaluation (Phases 1–4)

The NAIRR system should be designed to achieve its objective and goals in a deliberate manner. A “theory of change” for the NAIRR—that is, a causal model or map of how the goals of a program are intended to be achieved—can inform this process and provide a framework for its planning and evaluation. This includes articulating the inputs (i.e., available resources to leverage), activities (i.e., actions or work conducted to advance the program), outputs (i.e., the immediate, practical benefits of the program), outcomes (i.e., medium-term results), and longer-term impacts of the overall NAIRR effort, and how each successively feeds into the next (see Figure 6 for illustrative examples). The NAIRR is envisioned as a complex system with numerous entities responsible for creating, operating, and overseeing its components, integration, services, and policies.




Figure 6. Example Elements of a Theory of Change for the NAIRR

NAIRR governance entities should adopt a standard evaluation framework predicated on a clearly defined theory of change. The Steering Committee and Operating Entity, in collaboration with other NAIRR entities, should develop and publish appropriate KPIs based on this framework during NAIRR Implementation Phases 1–2, and adapt them as needed as the system matures. KPIs should be developed early with input from experts in program evaluation to ensure that data-collection mechanisms are built into NAIRR processes in a timely and reproducible manner (i.e., specific, measurable, attributable, realistic, and targeted).

To ensure objective and rigorous evaluation of the Operating Entity, resource providers, and overall NAIRR performance, the governance entities must enlist an expert and independent entity to act as its evaluator. Evaluation should be conducted against appropriate baseline measures and “counterfactuals” (i.e., scenarios or proxies for a particular outcome or metric that would occur in the absence of the NAIRR).

To assess the performance of the NAIRR system and its progress toward achieving its four goals, its cognizant entities must plan for and participate in periodic, independent evaluation. Ideally, the NAIRR system should be designed and established deliberately, using appropriate inputs to its activities for achieving near-term outputs, longer-term outcomes, and high-level impacts. Evaluation should be conducted at the level of (1) all aspects of the NAIRR system as a whole, (2) the Operating Entity, (3) the resource providers, and (4) individual research projects and users making use of the NAIRR. All four assessments may be conducted via one evaluation process conducted by the external evaluator. The NAIRR Steering Committee should develop KPIs for each entity in collaboration with NAIRR constituents and in alignment with NAIRR goals. KPIs can be technical, such as total computational power; usage-related, such as access counts for datasets or training tools; or human-centered, such as number of users. The NAIRR should be architected to facilitate the capture of KPIs that can be readily accessed through a dashboard and made available to the Steering Committee, Program Management Office, and Operating Entity. KPIs should also address diversity and equity—for example, not only the number of users, but also the demographics and institutional diversity of users.

The Operating Entity should develop clear expectations for each resource provider, including milestones and deliverables, tied to the KPIs and consistent with the mission of the NAIRR. The expectations should be reviewed by the Program Management Office, the Steering Committee, the User Committee, and the advisory boards and posted on the NAIRR website. There should be a mid-term evaluation of each resource provider by an external evaluator selected by the Operating Entity and approved by the Program Management Office. Failure of a resource provider to perform according to expectations should trigger a probationary period. Continued or longer-term failure to perform should result in decommissioning a resource provider.

The KPIs for the NAIRR resource providers and Operating Entity should be a limited set of high-level metrics that the Program Management Office can initially use to monitor and evaluate the operational effectiveness of the research resources coordinated and the services provided by the Operating Entity to the user community. These metrics should be clearly stated and published. KPIs should be vetted by the Steering Committee, the Program Management Office, and User Committee, and published for public comment. These metrics should form the basis of the RFPs for resource providers and for subsequent program calls. Responsibility for defining KPIs for key NAIRR units is summarized in Table 3.

Table 3. Roles in KPI Definition and Frequency of External Evaluation

NAIRR Unit to Be Evaluated	KPIs Defined by	Frequency of Reporting by External Evaluator
Overall NAIRR performance	Steering Committee, with input from Program Management Office, Operating Entity, User Committee, advisory boards	Annual

NAIRR Unit to Be Evaluated	KPIs Defined by	Frequency of Reporting by External Evaluator
Operating Entity	Steering Committee, with input from Program Management Office, Operating Entity, User Committee, advisory boards	Annual
Resource Providers	Operating Entity, with input from User Committee and advisory boards	Mid-term and Annual

Evaluation activities include planning and preparation, information gathering and assessment, and release of and response to evaluator findings. Since the NAIRR requires substantial startup time, the evaluation itself should be phased in over several years. In Phases 2–3, for example, the evaluation should focus on NAIRR’s inputs, activities, and outputs, which would be primarily process-driven, and on establishing baselines for longer-term outputs and outcomes. These initial evaluations should focus on implementation by the NAIRR and the Operating Entity as well as the resource providers. Subsequent evaluations should begin to evaluate progress towards the intended goals and outcomes of the NAIRR itself. The evaluation should expand in years 4–6 to include outcomes, while years 7–9 should also evaluate and measure progress toward the broader impacts.

Roadmap for Implementation

An infrastructure as complex as the NAIRR would require several years before it is fully operational, although the NAIRR is expected to reach its initial operational capability, in which it can begin to serve its envisioned user base, approximately two years after program initiation. Detailed implementation steps for key actors are summarized in Figure 7 for the four phases defined for establishment of the NAIRR: (1) Program Initiation and Operating Entity Selection, (2) Operating Entity Startup, (3) NAIRR Initial Operational Capabilities, and (4) NAIRR Steady-State Operations.

NAIRR Implementation Roadmap




Figure 7. NAIRR Implementation Roadmap

Steps to Initiate the NAIRR in 2023: Actions for the U. S. Government

Congress and the Federal agencies should take the following actions in 2023, as part of Phase 1: NAIRR Program Initiation, to begin establishing the NAIRR.

For the President and Executive Branch Departments and Agencies

The development and sustainment of the NAIRR will require active involvement by many Federal agencies, which will need to participate in the Steering Committee and the Program Management Office, allocate funds for the resource providers, and oversee the NAIRR's execution. The agency serving as the administrative home will need to establish a Program Management Office and allocate funds for the Operating Entity.

For the NAIRR to be successful, it will need to reach all major AI-using research communities—and for that to occur, all of the Federal research agencies that invest in AI R&D will need to participate in the management and funding of the NAIRR.

For Congress

Congressional legislation has continually reaffirmed the Federal Government's commitment to funding cutting-edge information technology R&D. The success of the NAIRR initiative will depend on similar commitments from the Federal Government using similar legislative tools and authorities. The long-term continuation of U.S. strategic advancement and leadership in AI depends on guidance and commitment from Congress. (See Appendix I for proposed NAIRR authorizing legislation drafted by the Task Force.)

6. Conclusion

AI is an engine of innovation that is already driving scientific discovery and economic growth, and is an integral component of solutions that stand to impact everything from routine daily tasks to societal-level challenges. To realize this promise, we must provide opportunities for researchers throughout the Nation to pursue cutting-edge AI research. As a Nation, we must come together to expand access to the resources that fuel AI, providing pathways for more Americans to pursue AI R&D and to access state-of-the-art resources. The NAIRR can help to broaden the range of researchers involved in AI, growing and diversifying approaches to and applications of AI. The NAIRR can help create opportunities for progress across all scientific fields and disciplines, including in critical areas such as AI auditing, testing, and evaluation; trustworthy AI; bias mitigation; and AI safety. Increased access and diversity of perspectives would, in turn, lead to new ideas that would not otherwise materialize and set the conditions for developing AI systems that are inclusive by design. The vision for a NAIRR laid out in this final report of the NAIRR Task Force can help meet this national need through a shared research cyberinfrastructure connecting researchers to the resources and tools that fuel AI R&D. The Task Force has presented a roadmap for doing so in a manner that builds from existing Federal investments; designs-in protections for privacy, civil rights, and civil liberties; and promotes diversity and equitable access. The NAIRR can help transform the U.S. national AI research ecosystem by strengthening and democratizing foundational, use-inspired, and translational AI R&D in the United States.

Appendix A. Definitions

Artificial intelligence (AI): See Box 1.

Cyberinfrastructure: Refers to infrastructure based on distributed computer, information, and communication technologies, including the enabling hardware, algorithms, software and services, communications, institutions, and expertise.

Experimental System/Resource: A system or resource that is exploring a new hardware or software capability and may provide an immature or rapidly evolving environment for the user to run in. Users may expect additional efforts to port applications to properly use the capabilities of the system, rather than a “turnkey” environment, and not all use cases may be well-supported.

Federated system: A set of semi-autonomous, decentralized resources that use a standard set of protocols allowing for integration, interoperability, and information sharing.

Foundational AI research: Discovery-oriented fundamental research that seeks to advance the frontiers of AI, including knowledge representation, reasoning, planning, learning, language processing, perception, vision, motion and manipulation, and so on.

Fundamental Research: Also known as basic research; spans the full spectrum from foundational, discovery-oriented to use-inspired, solution-oriented research.

National AI Research Resource (NAIRR): See Box 2.

On-premise: Computational hardware that is physically located on the premises of the organization making use of it, in contrast to remote hardware such as a commercial cloud.

Research on AI: Foundational, use-inspired, and translational research that advances scientific understanding of the nature of intelligence, mathematical understanding of the behavior of adaptive/autonomous systems, or algorithmic understanding of techniques in the component areas of AI (which include perception, learning, planning, and robotics) as well as research related to robustness, scalability, reliability, safety, security, privacy, interpretability, and equity of AI systems.

Testbeds: Platforms used to conduct research and validate theories, tools, or technologies in a rigorous, replicable manner. AI testbeds may take the form of simulated, live, or blended environments that support prototyping, development, and testing of AI applications that are robust and trustworthy. The concept of a testbed can encompass the environment itself—hardware and software—as well as the datasets and frameworks that support evaluation and the talent needed to manage the resource. AI testbeds may take the form of comparison testbeds (allowing researchers to measure the effectiveness of new engineering, math, or algorithmic developments) or validation testbeds (allowing developers to decide whether an end-to-end system is acceptable to move up the maturity cycle to a more advanced phase of development).

Translational AI research: Research that bridges foundational and use-inspired research with the delivery and deployment of its outcomes to the target community, and that supports essential bi-direction interplays where the delivery and deployment process informs the research; as in, translating research results from the lab to the market and society.

Use-inspired AI research: Fundamental research in AI that is motivated or inspired by particular use cases, and seeks to advance both the frontiers of AI and the specific use cases.

Appendix B. Details of NAIRR Task Force Establishment and Approach to Roadmap Development

Charge to the NAIRR Task Force

Congress charged the Task Force with proposing a national solution to provide researchers and students across scientific fields and disciplines with access to data and computing resources for AI R&D, along with appropriate educational tools and user support. Specifically, Congress directed the Task Force to develop a roadmap and implementation plan for establishing the NAIRR. The Task Force was launched on June 10, 2021, as a Federal Advisory Committee co-chaired by the National Science Foundation and the White House Office of Science and Technology Policy, and includes representatives from the U.S. Government, academia, and the private sector. Its members' expertise spans foundational, use-inspired, and trustworthy AI R&D, as well as research cyberinfrastructure. This report constitutes the Task Force's final deliverable, pursuant to its Congressional mandate. Congress specified that the NAIRR roadmap and implementation plan address nine key dimensions, as stated in Box B.1. The Task Force activities were bounded to developing recommendations and proposing a roadmap and implementation plan for a NAIRR to the President and to Congress. The Task Force will conclude its work within 90 days after submission of this final report; the Task Force itself will not execute any of its recommendations, nor will it be involved in the administration of a future NAIRR.

Box B.1. Required Elements of the NAIRR Roadmap and Implementation Plan³

- (1) IN GENERAL—The Task Force shall develop a coordinated roadmap and implementation plan for creating and sustaining a National Artificial Intelligence Research Resource.
- (2) CONTENTS—The roadmap and plan required by paragraph (1) shall include the following:
 - A. Goals for establishment and sustainment of a National Artificial Intelligence Research Resource, and metrics for success.
 - B. A plan for ownership and administration of the National Artificial Intelligence Research Resource, including i. an appropriate agency or organization responsible for the implementation, deployment, and administration of the Resource; and ii. a governance structure for the Resource, including oversight and decision-making authorities.
 - C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.
 - D. Capabilities required to create and maintain a shared computing infrastructure to facilitate access to computing resources for researchers across the country, including scalability, secured access control, resident data engineering and curation expertise, provision of curated datasets, computational resources, educational tools and services, and a user interface portal.
 - E. An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government datasets as part of the National Artificial Intelligence Research Resource.
 - F. An assessment of security requirements associated with the National Artificial Intelligence Research Resource and its research and a recommendation for a framework for the management of access controls.
 - G. An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.
 - H. A plan for sustaining the Resource, including through Federal funding and partnerships with the private sector.
 - I. Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities and milestones to implement the Resource.

Task Force Approach

The Task Force’s work was divided into phases. The first phase began with the Task Force’s first convening in July 2021 and culminated with the release of the interim NAIRR report in May 2022. The second phase, between May 2022 and January 2023, was devoted to the development of the final report and roadmap.

Initial Phase

During the initial phase, the Task Force convened seven virtual public meetings to discuss and deliberate on key NAIRR uses, potential impacts, system requirements, and design elements. At these meetings, the Task Force heard from expert briefers and panelists to augment the members’ own expertise, and to ensure that multiple perspectives and experiences were considered in Task Force discussions and deliberations. A complete list of invited panelists as well as respondents to the first request for information (RFI), published in July 2021, can be found in Appendix E.

The NAIRR Task Force submitted its interim report to the President and Congress in May 2022.¹⁴ It set forth the Task Force’s vision for the NAIRR, along with preliminary recommendations on the nine key areas identified by Congress (see Box B.1).

Final Phase

During the final phase, the Task Force convened four virtual public meetings to discuss and deliberate on how best to implement the plan published in the interim report. At these meetings, the Task Force heard from several expert briefers and panelists. Topics addressed by these invited experts included international perspectives and associated Federal efforts for the provisioning of data and computing.

The Task Force also reviewed 23 public responses to a May 2022 RFI asking for comment on the interim report and potential approaches to implementation. These responses reflect feedback from individuals (ranging from academics to interested members of the public), groups, and organizations (spanning non-profits, civil society groups, research organizations, and small and large businesses). For a full list of respondents and a link to these RFI responses, see Appendix C. In the course of their deliberations during this phase, Task Force members also engaged with additional outside subject matter experts (see Appendix D for a complete list of experts consulted) in support of their considerations toward this final report. A public listening session was held on June 23, 2022, to provide another opportunity for the public to provide input. Seventy-four individuals registered to participate in the listening session, of whom 48 attended. Eight of those individuals spoke at the meeting, including three from civil society or advocacy groups, one from academia, one from an industry or industry association group, one from government, one private citizen, and one other. See Appendix E for a complete list of participants and speakers at this session.

Appendix C. Briefers to the Task Force

The Task Force held eleven public meetings between its launch in July 2021 and the release of this final report. At these meetings, the Task Force discussed and developed a vision for the NAIRR, heard input from invited expert speakers and panelists, and deliberated on key findings and preliminary recommendations for the design of the NAIRR and its roadmap and implementation. These outside expert briefers and panelists, along with their affiliations, are listed here.

July 28, 2021

The STRIDES program

Andrea Norris & Nick Weber, *National Institutes of Health*

August 20, 2021

Value proposition and intended outcomes of a NAIRR

Damian Clarke, *Chief Information Officer and Computer Science Faculty, Alabama A&M University*

James Deaton, *Executive Director, Great Plains Network*

Deborah Dent, *Chief Information Officer, Jackson State University*

Tripti Sinha, *Assistant Vice President and Chief Technology Officer, University of Maryland, and Executive Director of the Mid-Atlantic Crossroads (MAX)*

Talitha Washington, *Director, Atlanta University Center Consortium Data Science Initiative*

Ownership, governance, and administration models

Sharon Broude Geva, *Director for Innovation and Computational Research, University of Michigan*

Manish Parashar, *Office Director, Office of Advanced Cyberinfrastructure, National Science Foundation⁵⁶*

Gina Tourassi, *Director, National Center of Computational Sciences and the Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory*

John Towns, *Executive Associate Director for Engagement, National Center for Supercomputing Applications and Deputy CIO for Research IT, University of Illinois at Urbana-Champaign*

Frank Würthwein, *Interim Executive Director, San Diego Supercomputer Center*

October 25, 2021

Data resources

Ian Foster, *Director, Data Science and Learning Division, Argonne National Laboratory; Professor of Computer Science, University of Chicago*

Robert L. Grossman, *Professor of Medicine and Computer Science, University of Chicago*

Ron Hutchins, *Vice Provost for Academic Technologies, University of Virginia*

Anita Nikolich, *Research Scientist and Director of Research and Technology Innovation, University of Illinois at Urbana-Champaign*

Nancy Potok, *CEO, NAPx Consulting; former Chief Statistician of the United States*

Andrew Trask, *Leader, OpenMined*

User resources: portal interface, educational tools

Tiziana Ferrari, *Director, EGI Foundation*

Kimberly Greene Starks, *Global Lead, Infrastructure and Technology Strategy, IBM University Programs*

Ana Hunsinger, *Vice President for Community Engagement, Internet2*

Ed Lazowska, *Professor and Bill & Melinda Gates Chair Emeritus, Paul G. Allen School of Computer Science & Engineering, University of Washington*

December 13, 2021

Privacy, civil rights, and civil liberties requirements

Solon Barocas, *Principal Researcher, Microsoft Research; Adjunct Assistant Professor, Information Science, Cornell University*

Lujo Bauer, *Professor, Electrical & Computer Engineering and Computer Science, Carnegie Mellon University*

danah boyd, *Partner Researcher, Microsoft Research; and Founder/President, Data & Society*

Deborah Raji, *Fellow, Mozilla Foundation*

Nicol Turner Lee, *Senior Fellow and Director of the Center for Technology Innovation, Brookings Institution*

Hannah Quay-de la Vallee, *Senior Technologist, Center for Democracy and Technology*

February 16, 2022

User perspectives on the NAIRR

Tom Dietterich, *Distinguished Professor Emeritus in Computer Science, Oregon State University*

Susanta Ghosh, *Assistant Professor in Mechanical Engineering-Engineering Mechanics, Michigan Technological University*

Kinnis Gosha, *Hortinius I. Chenault Endowed Associate Professor of Computer Science, Morehouse College*

Gail Rosen, *Professor, Drexel University*

Rima Seilova-Olson, *Co-Founder and Chief AI Scientist, Kintsugi*

Carlos Theran, *Research Associate, Florida A&M University*

April 8, 2022

Building responsible AI review processes for the NAIRR

Beena Ammanath, *Author, Trustworthy AI and Head of Global Deloitte AI Institute*

Michael Bernstein, *Associate Professor of Computer Science, Stanford University*

Arvind Narayanan, *Associate Professor of Computer Science, Princeton University*

Beth Plale, *Professor and Director of the Data to Insight Center, Indiana University Bloomington*

Christo Wilson, *Associate Professor of Computer Science, Northeastern University*

May 20, 2022

No external speakers; the only agenda item was for the Task Force to vote on the interim report.

July 25, 2022

International perspectives on the NAIRR

Karine Perset, *Head, AI Unit, Division for Digital Economy Policy, OECD*

Mark Leggott, *Director of International Relations, Digital Research Alliance of Canada*

Renaud Vedel, *Chief of Staff to the Minister for the Digital Economy, France*

Kazuyuki Takada, *Director, Industrial Science and Technology Project Promotion Office, Ministry of Economy, Trade and Industry (METI), Japan*

Alison Kennedy, *Strategic Adviser, Science and Technology Facilities Council, UK Research and Innovation*

Eliana Cardoso Emediato de Azambuja, *General Coordinator of Digital Transformation, Department of Science, Technology and Digital Innovation, Secretariat of Entrepreneurship and Innovation, Ministry of Science, Technology and Innovation, Brazil*

September 12, 2022

Associated Federal efforts for provision of data and computing

Shelly Martinez, *Senior Statistician, Office of Management and Budget*

Vipin Arora, *Deputy Director, National Center for Science and Engineering Statistics,
National Science Foundation*

Kamie Roberts, *Director, National Coordination Office for the Networking and Information
Technology Research and Development Program*

Jerry Sheehan, *Deputy Director for Policy and External Affairs, National Library of
Medicine, National Institutes of Health*

October 21, 2022

No external speakers; the only agenda item was for the Task Force to deliberate on the final report.

January 13, 2023

No external speakers; the only agenda item was for the Task Force to vote on the final report.

Appendix D. Public Input Provided on the Interim Report in Response to the Federal Request for Information

Concurrently with the publication of the interim report, the Task Force issued a Request for Information (RFI) to solicit public feedback on the Task Force’s preliminary findings and recommendations outlined in the interim report, and particularly, on how the recommendations could be successfully implemented. The RFI was open for comments from May 25, 2022, through June 30, 2022. This RFI received 24 responses. The list of respondents to this RFI follows. The full texts of the responses are available at <https://www.ai.gov/87-fr-31914-responses/>.

- ACT | The App Association
- American Psychological Association (APA)
- Anthropic
- Centre for the Governance of AI (GovAI)
- Consumer Reports
- Data Foundation
- Dreifus, Greg and Caso, Luis Videgaray
- Electronic Privacy Information Center (EPIC)
- Engine
- Hugging Face
- IBM
- IEEE – USA
- Internet2
- SeedAI
- Shavit, Yonadav; Kaushik, Divyansh; Lipton, Zachary C.; Bowman, Samuel R.; and Goldner, Kira
- Sheehan, Matt; Critch, Andrew; Jackson, Krystal; and Feldgoise, Jacob
- Software & Information Industry Association (SIIA)
- Stanford Institute for Human-Centered Artificial Intelligence (HAI)
- The MITRE Corporation
- U.S. Chamber of Commerce Technology Engagement Center
- University of Arizona, CODATA Center of Excellence in Data for Society
- University of Southern California (USC) Information Sciences Institute (ISI)
- Wehbe, Joseph
- Wieder, Robin

Appendix E. Public Input Provided on the Initial Federal Request for Information on Designing the NAIRR

A Request for Information on the design of a NAIRR was posted in the Federal Register (86 FR 39081) on July 23, 2021; the comment period closed on October 1, 2021. The Task Force received 84 responses. The list of respondents to this Request for Information follows; the full text of the responses is available at <https://www.ai.gov/nairrf/86-fr-39081-responses/>.

- Abdoli, Abas; Coffee, Ryan N.; Edelen, Auralee; Kagan, Michael; Ratner, Daniel; Reddy, Sohail; and Terao, Kazuhiro
- Accenture
- ACM U.S. Technology Policy Committee
- The Aerospace Corporation
- AI Now Institute of New York University and Data & Society Research Institute
- AI Redefined, Inc.
- The Alexandria Archive Institute (Open Context)
- Amazon Web Services
- American Civil Liberties Union (ACLU)
- American Psychological Association (APA)
- Anthropic
- Argonne National Laboratory
- Atlantic Council GeoTech Center
- August, Michael
- BeeHero
- Booz Allen Hamilton
- C-2
- Cadence
- CalypsoAI Corp.
- Carnegie Mellon University
- Center for Data Innovation
- Center for Democracy and Technology
- Center for Security and Emerging Technology
- Cerner Corporation
- Computing Community Consortium, Computing Research Association-Industry, and the Association for the Advancement of Artificial Intelligence
- Consumer Reports
- CrowdAI
- Deloitte
- Digital Diagnostics
- Domalpally, Amitha and Channa, Roomasa
- Ekins, Sean
- Electronic Privacy Information Center (EPIC)
- Engine
- The Enterprise Neurosystem
- FABRIC Testbed
- Feddema, John T.; Stracuzzi, David J.; and Steward, James R.
- Freed, Ben and Choset, Howie
- Freeman, Jared; Leins, Drew; and Gaffney, Niall
- Ghosh, Aishik
- Gilmore, Wayne; Goodhue, John; Hill, Christopher N.; Kaelli, David; Kolaczyk, Eric; Kurose, Jim; and Yackel, Scott
- Google
- Hewlett Packard Enterprise
- Hyperion Research
- IBM
- Indiana University
- Infiltron
- Information Technology Industry Council

- Institute of Electrical and Electronics Engineers (IEEE) Standards Association
- Internet2
- Kapoor, Savash; Kshirsagar, Mihir; and Narayanan, Arvind
- Kubitz, Kermit
- Lawrence Berkeley National Laboratory
- Lawrence Berkeley National Laboratory Machine Learning Group
- Lawrence Livermore National Laboratory
- Mathematica
- Medical Imaging and Resource Center, University of Chicago
- Microsoft
- The MITRE Corporation
- Moffitt Cancer Center
- NASA
- NSF AI Institute for Artificial Intelligence and Fundamental Interactions
- NSF AI Institutes
- NVIDIA
- National Center for Atmospheric Research
- National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign
- National Energy Technology Laboratory
- NiyamIT, Inc.
- Noblis
- Northeastern University
- Open Commons Consortium at the Center for Computational Science Research, Inc.
- Oracle America, Inc.
- Ossorio, Pila
- Palantir Technologies, Inc.
- Partnership on AI
- Patterson, Maria
- Representatives from the National Oceanic and Atmospheric Administration (NOAA) Artificial Intelligence Executive Committee (NAIEC) and the Center for Artificial Intelligence (NCAI)
- SAS
- Sirintrapun, Joseph S.
- Stanford Libraries
- Stanford University Institute for Human-Centered Artificial Intelligence (HAI)
- U.S. Chamber of Commerce Technology Engagement Center
- University of Florida
- University of Illinois, Chicago
- Xiao, Steve
- Yankeelov, Thomas

Appendix F. Subject Matter Experts Consulted by Task Force Members

Pete Beckman
Argonne National Laboratory

Jim Brase
*COVID-19 HPC Consortium
and Lawrence Livermore National
Laboratory*

Sandeep Chandra
San Diego Supercomputer Center

Kate Crawford
AINow (NYU)

Ian Ferreira
Core Scientific, Inc.

Brett Goldstein
Vanderbilt University

Julie Haney
*National Institute
of Standards and Technology*

Nick Hart
Data Foundation

Robert Jackson
Spherecom Enterprises

Suzette Kent
Kent Advisory Services

Christine Kirkpatrick
San Diego Supercomputer Center

Tony LaVoi
*National Oceanic
and Atmospheric Administration*

Aaminah Norris
Algorithmic Justice League

Jason Owen-Smith
University of Michigan

Joris Poort
ReScale, Inc.

Nancy Potok
NAPx Consulting

Catherine Schuman
University of Tennessee, Knoxville

Adam Schwartz
Ames Laboratory

Brock Webb
U.S. Census Bureau

Harlan Yu
Upturn

Appendix G. NAIRR Public Listening Session

On Thursday, June 23, 2022, the Task Force hosted a listening session to collect public input on the initial findings and recommendations of the interim report. A notice in the Federal Register announcing the session was released on May 25, 2022. The notice is available at <https://www.federalregister.gov/documents/2022/05/25/2022-11222/public-listening-session-on-implementing-initial-findings-and-recommendations-of-the-national>.

The 74 registrants for the session indicated affiliation with academia, civil society or advocacy groups, government, industry or industry association groups, private citizens, and others. Of these registrants, 13 indicated a desire to speak.

Science and Technology Policy Institute researchers opened the meeting by introducing the agenda and goals of the session. NAIRR Task Force Co-Chairs Dr. Manish Parashar and Dr. Lynne Parker then provided a short briefing to participants on the NAIRR Task Force's work to provide context in advance of public comments. During the session, 48 individuals attended, and there were eight speakers, including three from civil society or advocacy groups, one from academia, one from an industry or industry association group, one from government, one private citizen, and one other. The session had been scheduled to last for two hours with the possibility of ending early if all interested speakers had been heard. With a limited number of participants interested in speaking, the session lasted about 55 minutes and ended early.

Appendix H. NAIRR Task Force Staff and Contributors

Emily Grumbling

IDA Science and Technology Policy Institute

Matthew Christman

IDA Science and Technology Policy Institute

Matthew Ishimaru

IDA Science and Technology Policy Institute

Morgan Livingston

IDA Science and Technology Policy Institute

Logan Practico

IDA Science and Technology Policy Institute

Michelle Tolbert

*Networking and Information Technology
Research and Development Program*

Lisa Van Pay

IDA Science and Technology Policy Institute

Taylor White

IDA Science and Technology Policy Institute

Brian Zuckerman

IDA Science and Technology Policy Institute

Kevin Garrison

Institute for Defense Analyses

Patricia Sadiq

Institute for Defense Analyses

Geoff Holdridge

*National Nanotechnology Coordination
Office*

Appendix I. Examples of NAIRR Evaluation Metrics

Data and evidence will be needed for both monitoring and evaluation of the NAIRR system and activities. A budget must be established for the external evaluator and paid for from Operating Entity funds. The Operating Entity and Program Management Office, with input from NAIRR advisory boards and the NAIRR Steering Committee, must agree to a theory of change for designing the NAIRR activities and infrastructure that will serve as the basis for evaluation. To ensure rigor and objectivity, the evaluation should be conducted by an independent, external entity with expertise in program evaluation.

The following tables provide examples of potential evaluation metrics that might be associated with the inputs, activities, outputs, and outcomes for a NAIRR theory of change (Table H.1), as well as for measuring progress toward the four goals of the NAIRR (Table H.2). Where possible, metrics should be automatically collected and made available in real time to the entities involved in NAIRR governance. All metrics should be assessed relative to a counterfactual to the extent possible. This could include pre-NAIRR baseline metrics and associated projections, or metrics for an analogous discipline or research community that has not had the same intervention (that is, does not have a dedicated, federally funded, R&D cyberinfrastructure) over the same period.

In addition to overall NAIRR performance, the Operating Entity and individual resource providers must be evaluated. Resource providers should be evaluated for operational efficiency on the following high-level performance metrics: Customer support, queue times, consultant response time, computational time and services, allocated time limits, and quality and completeness of resource documentation. The characteristics of provisioned resources and associated needs may vary, including by user community. Additional specific metrics for each major category of provisioned resources should complement the high-level performance metrics. The overall portfolio of research supported via the NAIRR should also be evaluated as part of NAIRR evaluation to support strategic adjustments.

Table H.1. Examples of Metrics that Could be Associated with a NAIRR Theory of Change

Input Metrics	Activity Metrics	Output Metrics	Outcome Metrics
<ul style="list-style-type: none"> • Number and type of computational and data resources leveraged • Amount of funding from Federal agencies • Amount of funding and in-kind support provided by philanthropic organizations • Amount of funding and in-kind support provided by industry • Staff time (in full-time employment equivalents) • Expertise included among staff 	<ul style="list-style-type: none"> • Number of RFPs drafted for NAIRR resources • Number of cross-agency NAIRR competitions launched; proposal and acceptance rates • Number of and variety of workshops held • Frequency and extent of outreach activities • Amount of funding allocated to each resource/service type (compute, data, user training and support, testbeds) • Number and diversity of individuals working on research conducted on the NAIRR 	<ul style="list-style-type: none"> • Number and variety of resources available to users via NAIRR • Computational capacity available for allocation via NAIRR • Number of high-quality data sets available • Number of key information and training resources available over time • Resource access statistics, including processor hours allocated • Consistency of resource availability • Number and diversity of new NAIRR users • Number and diversity of NAIRR users newly engaging with AI • Number and diversity of NAIRR-mediated collaborations • Number and diversity of users leveraging training materials 	<ul style="list-style-type: none"> • Number and diversity of NAIRR users working in academia, the private sector, and non-profits • Earnings and employment outcomes of NAIRR users working in academia, the private sector, and non-profits • Number of startups established by NAIRR users • Productivity and growth of firms associated with NAIRR users or as vendors (including new vendor startups) to research conducted on the NAIRR

Table H.2. Examples of Metrics for Assessing Progress toward NAIRR Goals

Innovation	Diversity	Capacity	Trustworthy AI
<ul style="list-style-type: none"> • Number of startups established by NAIRR users • Number of startups emerging from research conducted on the NAIRR • Productivity and growth of firms associated with NAIRR users • Productivity and growth of vendors (including new vendor startups established) to support NAIRR resource/service providers • Number of “groundbreaking” publications and patents across S&E that can be traced to NAIRR users 	<ul style="list-style-type: none"> • Number and share of AI “research-involved” individuals from underrepresented or underserved populations • Earnings and employment outcomes of AI “research-involved” individuals from underrepresented or underserved populations when placed • Institutional demographics of AI researchers and NAIRR users • Demographics of NAIRR users, Operating Entity leaders, and governance entities 	<ul style="list-style-type: none"> • Number of AI “research involved” individuals (defined as individuals paid on AI grants or in AI jobs) • Earnings of AI “research involved” individuals when placed • Number of individuals leveraging NAIRR for education and training • Number of AI-intensive firms with establishment linkable to NAIRR • Employment in AI-intensive firms with establishment linkable to research conducted on the NAIRR • Number of research publications, patents, and awards in AI and at the intersection of AI and other fields traceable to NAIRR users 	<ul style="list-style-type: none"> • Number of tools developed for trustworthy AI leveraging NAIRR • Access statistics for NAIRR’s AI ethics education and training tools • Number and impact of papers published on AI ethics/trustworthy AI and citing NAIRR • Share of AI publications that address AI ethics, trustworthiness, and societal implications • Number, use statistics, and efficacy of NAIRR ethics tools and trainings • Extent of NAIRR engagement with AI ethics experts • Representation of social science and AI ethics expertise in NAIRR governance entities • Expenditures on tools, trainings, services, and consultations related to AI ethics

Appendix J. Draft Legislative Language for NAIRR Authorization

The following text represents the NAIRR Task Force’s best efforts to capture its recommendations in legislative text, with annotations to explain the intent of the Task Force.

SECTION 1. SHORT TITLE.

This Act may be cited as the “National Artificial Intelligence Research Resource Act” or the “NAIRR Act.”

SEC. 2. NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE.

The National Artificial Intelligence Initiative Act of 2020 (15 U.S.C. 9411 et seq.)¹ is amended by adding at the end the following:

“TITLE LVI—NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE

“SEC. 5601. FINDINGS.

“The Congress finds the following:

“(1) Much of today’s cutting-edge artificial intelligence research relies on access to computational resources and large datasets.

“(2) Access to the computational resources and datasets necessary for artificial intelligence research and development is often limited to very large technology companies and well-resourced universities.

“(3) The lack of access to computational and data resources has resulted in insufficient diversity in the artificial intelligence research and development community.

¹ The National Artificial Intelligence Initiative Act of 2020 (NAIIA) appears as division E of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021 (Pub. L. 116–283) (FY21 NDAA). This draft legislation does not include potential technical conforming amendments (e.g., to the table of sections in the NAIIA or the FY21 NDAA) necessary to execute the substantive amendment recommended.

“(4) Engaging the full and diverse talent of the United States is critical for maintaining United States leadership in artificial intelligence and ensuring that artificial intelligence is developed in a manner that benefits all Americans.

“(5) The National Artificial Intelligence Research Resource Task Force, authorized under section 5106, recommended the establishment of a National Artificial Intelligence Research Resource in a report entitled “Strengthening and Democratizing the U.S. Artificial Intelligence Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource” on January 24, 2023.

“SEC. 5602. DEFINITIONS.²

“In this title:

“(1) ADVISORY BOARDS.—The term ‘Advisory Boards’ means the advisory boards established in section 5603(d).

“(2) AI TESTBED.—The term ‘AI testbed’ means a simulated, live, or blended environment that support prototyping, development, and testing of an artificial intelligence application, including—

“(A) the hardware or software for the environment required for an artificial intelligence application;

“(B) data sets and frameworks that support evaluation of an artificial intelligence application; and

“(C) the individuals required to manage an artificial intelligence application.

“(3) ETHICS ADVISORY BOARD.—The term ‘Ethics Advisory Board’ means the advisory board described in section 5603(d)(2)(C).

“(4) EXECUTIVE AGENCY.—The term ‘Executive agency’ has the meaning given such term in section 105 of title 5, United States Code.

“(5) NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH RESOURCE AND NAIRR.—The terms ‘National Artificial Intelligence Research Resource’ and ‘NAIRR’ have the meaning given the term ‘National Artificial Intelligence Research Resource’ in section 5106(g).

“(6) NATIONAL SECURE DATA SERVICE.—The term ‘National Secure Data Service’ means the demonstration project established in section 10375 of the Research and

² Because the new title LVI is added to the NAIIA, the definitions located in section 5002 of NAIIA ([15 USC 9401](#)) that apply across the entirety of the NAIIA apply in this title as well without explicit reference in this new title. Those defined terms include: (1) Advisory Committee; (2) agency head; (3) artificial intelligence; (4) community college; (5) Initiative; (6) Initiative Office; (7) Institute; (8) institution of higher education; (9) Interagency Committee; (10) K-12 education; and (11) machine learning, though not all of those terms are used in this new title.

Development, Competition, and Innovation Act³ (42 U.S.C. 19085) or any successor program.

“(7) OPEN SOURCE SOFTWARE.—The term ‘open source software’ has the meaning given such term in section 2201 of the Homeland Security Act of 2002 (6 U.S.C. 651).

“(8) OPERATING ENTITY.—The term ‘Operating Entity’ means the Operating Entity selected by the Program Management Office as described in section 5603(b)(4)(A).

“(9) PROGRAM MANAGEMENT OFFICE.—The term ‘Program Management Office’ means the Program Management Office established in section 5603(b).

“(10) RESEARCHER.— The term ‘researcher’ means a person who conducts research.

“(11) RESOURCE OF THE NAIRR.—The term ‘Resource of the NAIRR’ means a resource described in section 5604(b).

“(12) SCIENCE ADVISORY BOARD.— The term ‘Science Advisory Board’ means the advisory board described in section 5603(d)(2)(A).

“(13) STEERING COMMITTEE.—The term ‘Steering Committee’ means the committee described in section 5603(c).

“(14) STUDENT.— The term ‘student,’ when used with respect to an institution of higher education, means an individual who is—

“(A) registered as a student with the institution;

“(B) enrolled in not less than 1 class of the institution; or

“(C) otherwise considered a student in good standing by the institution.

“(15) TECHNOLOGY ADVISORY BOARD.—The term ‘Technology Advisory Board’ means the advisory board described in section 5603(d)(2)(B).

“(16) USER COMMITTEE.—The term ‘User Committee’ means the advisory board established in section 5603(d)(2)(D).

“SEC. 5603. ESTABLISHMENT; GOVERNANCE.

“(a) ESTABLISHMENT.—Not later than 12 months after the date of the enactment of the National Artificial Intelligence Research Resource Act, the Director of the National Science

³ Division B of what is commonly known as the CHIPS and Science Act (Pub. L. 117-167).

Foundation, in coordination with the Steering Committee, shall establish the National Artificial Intelligence Research Resource to—

- “(1) spur innovation in artificial intelligence research and development;
- “(2) increase diversity among researchers and students of artificial intelligence;
- “(3) improve capacity for artificial intelligence research in the United States; and
- “(4) advance the development of trustworthy artificial intelligence.

“(b) PROGRAM MANAGEMENT OFFICE.—

“(1) ESTABLISHMENT.—The Director of the National Science Foundation shall establish within the National Science Foundation a Program Management Office to oversee the day-to-day functions of NAIRR and shall appoint an individual, who may be from another Federal agency, to head the Program Management Office.

“(3) STAFF.—The head of the Program Management Office may identify staff and direct all employees of the Program Management Office, in accordance with the applicable provisions of title 5, United States Code.

“(4) DUTIES.—The duties of the Program Management Office shall include—

“(A) in coordination with the Steering Committee and Advisory Boards as appropriate—

“(i) developing the funding opportunity and solicit bids for the Operating Entity;

“(ii) selecting through a competitive and transparent process an organization to be designated the Operating Entity;

“(iii) overseeing the appointment of the Director and senior staff of the Operating Entity;

“(iv) overseeing compliance with the contractual obligations of the Operating Entity;

“(v) establishing evaluation criteria for the NAIRR;

“(vi) overseeing asset allocation and utilization;

“(vii) identifying an external independent evaluation entity; and

“(viii) assessing the performance of the Operating Entity on a periodic basis;
and

“(B) delegating, with appropriate oversight, operational tasks to the Operating Entity, including—

“(i) coordinating the provisioning of Resources of the NAIRR;

“(ii) maintaining a portal and associated services for users to access Resources of the NAIRR;

“(iii) developing NAIRR policies and procedures;

“(iv) hiring and managing a staff (including experts in cyber infrastructure management, data science, research design, privacy, ethics, civil rights and civil liberties, legal and policy matters) to support NAIRR operations;

“(v) continually modernizing NAIRR infrastructure;

“(vi) ensuring diversity, equity, inclusion, and accessibility in all aspects of the NAIRR, including operations;

“(vii) conducting ongoing evaluation and assessment of the NAIRR;

“(viii) establishing key performance indicators for the NAIRR, in coordination with the Steering Committee and Advisory Boards;

“(ix) publishing publicly-available annual reports reviewing the performance of the NAIRR, Resources of the NAIRR, and NAIRR governance structures;

“(x) establishing and administering training to new users on accessing a Resource of the NAIRR; research design; and issues related to privacy, ethics, civil rights and civil liberties, safety, and trustworthiness of artificial intelligence systems; and

“(xi) facilitating connections to AI testbeds.

“(c) STEERING COMMITTEE.—

“(1) ESTABLISHMENT AND MEMBERSHIP.—The Director of the Initiative Office shall establish a Steering Committee comprising agencies from the Interagency Committee as determined by the co-chairs of the Interagency Committee to have substantial expertise, have substantially funded or conducted artificial intelligence research and development, or have some other significant relationship with the NAIRR.

“(2) CHAIR AND CO-CHAIRS.—The Steering Committee shall be chaired by the Director of the Initiative Office. The Director of the Initiative Office may establish co-chairs of the Steering Committee based on members of the Steering Committee rotating on a pre-determined schedule.

“(3) CHANGES TO STEERING COMMITTEE COMPOSITION.—The Director of the Initiative Office shall review the composition of the Steering Committee and update the composition of the Steering Committee if necessary, not less frequently than every three years. A member of the Steering Committee may leave the Steering Committee as part of such a review.

“(4) SUBCOMMITTEES AND WORKING GROUPS.—

“(A) IN GENERAL.—The Steering Committee may establish subcommittees, working groups, or other permanent or temporary bodies of certain members of the Steering Committee.

“(B) WORKING GROUP ON COLLABORATING WITH THE FEDERAL INTERAGENCY COUNCIL ON STATISTICAL POLICY.—The Steering Committee shall establish a working group to assess options for establishing a secure node for the NAIRR to enable large-scale analysis of government data for statistical purposes in accordance with the Standard Application Process and, as practicable, as part of the National Secure Data Service.

“(5) DUTIES.—The Steering Committee shall—

“(A) coordinate with the National Science Foundation and the Program Management Office to oversee and approve the operating plan for NAIRR, request the budget for the NAIRR, develop and release a request for proposals to solicit bids for the Operating Entity, including establishing the terms and conditions and functions of the Operating Entity;

“(B) work with the Program Management Office to review candidates and select an entity to act as the Operating Entity;

“(C) identify resources that could be federated, participate in resource provider selection, and provide direction to the Operating Entity about resource allocation and how those resources should be made accessible via the NAIRR;

“(D) define key performance indicators for the NAIRR, in conjunction with the Program Management Office, User Committee, and Advisory Boards;

“(E) evaluate NAIRR performance against the key performance indicators defined in subparagraph (D) on a periodic basis and not less frequently than once every year;

“(F) develop an annual report transmitted to the Initiative Office and publicly released on the progress of the NAIRR that includes a summary of the evaluation concluded in subparagraph (E) and any recommendations for changes to NAIRR; and

“(G) oversee a periodic independent assessment of the NAIRR.

“(6) PROVISION OF RESOURCES OF THE NAIRR.—The agencies comprising the Steering Committee are authorized to provide the Operating Entity with a Resource of the NAIRR or funding for a Resource of the NAIRR.

“(d) ADVISORY BOARDS.—

“(1) IN GENERAL.—The head of the Program Management Office, acting through the Director of the Operating Entity, may establish Advisory Boards to provide advice to the Operating Entity and Program Management Office.

“(2) INITIAL ADVISORY BOARDS.—Not later than 3 months after the date of establishment of the NAIRR under subsection (a) the head of the Program Management Office, acting through the Director of the Operating Entity, shall establish the following Advisory Boards:

“(A) The Science Advisory Board comprising representatives from the scientific community, the public, public interest groups, the private sector, and other large-scale cyberinfrastructure projects to provide advice on the rapidly changing needs across multiple scientific domains.

“(B) The Technology Advisory Board comprising information technology experts from the private sector, government, and academia to provide advice on technological developments to aid the provisioning and use of Resources of the NAIRR and privacy and security technologies.

“(C) The Ethics Advisory Board comprising representatives of scientific societies, public interest groups, and government agencies to provide advice on ethics, fairness, bias, risks, privacy, civil rights, and civil liberties related to artificial intelligence.

“(D) The User Committee comprising representatives of different types of NAIRR users to provide recommendations on—

“(i) possible future directions for artificial intelligence research and training;

“(ii) user needs and requirements; and

“(iii) NAIRR policies and governance.

“(3) MEETING FREQUENCY.—Each Advisory Board shall meet not less frequently than twice per year.

“(4) COMPOSITION.—Each Advisory Board shall comprise members from government agencies, the private sector, academia, and public interest groups.

“(5) SELECTION.—The Director of the Operating Entity shall recommend individuals for the head of the Program Management Office to select from, after consultation with the Steering Committee.

“(6) TERMS.—Each member of an Advisory Board shall serve for a period of not more than three years. The terms of initial appointments to any Advisory Board may be staggered to allow for rotating members.

“(7) REPORTING.—The head of the Program Management Office shall, not less frequently than once per year, publicly report the following information for each Advisory Board:

“(A) Name of board.

“(B) Date of establishment.

“(C) Dates of meetings in the preceding 12 months.

“(D) Names and affiliations of members.

“(E) A list of formal reports or other documents produced and summaries of recommendations provided.

“(8) NONAPPLICABILITY OF FEDERAL ADVISORY COMMITTEE ACT.—The Federal Advisory Committee Act (5 U.S.C. App.) shall not apply to Advisory Boards.

“SEC. 5604. RESOURCES OF THE NAIRR.

“(a) IN GENERAL.—The head of the Program Management Office, acting through the Director of the Operating Entity and in coordination with the Steering Committee, Advisory Boards and User Committee, shall—

“(1) federate, coordinate, and allocate the provisioning of Resources of the NAIRR;

“(2) establish policies to govern the procurement and intake of Resources of the NAIRR;

“(3) establish policies on and review Resources of the NAIRR for concerns related to ethics, privacy, civil rights, and civil liberties, in coordination with the Ethics Advisory Board;

“(4) retire Resources of the NAIRR no longer available or needed; and

“(5) publicly report a summary of categories of available Resources of the NAIRR, categories of sources of such Resources of the NAIRR, and issues related to Resources of the NAIRR.

“(b) RESOURCES OF THE NAIRR.—The NAIRR shall offer at least the following resources:

“(1) A mix of computational resources, including—

“(A) on-premises, cloud-based, hybrid, and emergent resources;

“(B) not less than 1 large-scale machine-learning supercomputer;

“(C) public cloud providers providing access to popular computational and storage services for NAIRR users; and

“(D) specifying an open-source software environment for the NAIRR.

“(2) Data, including by—

“(A) publishing interoperability standards for data repositories and selecting and developing, through a competitive bidding process, repositories to be available to NAIRR users;

“(B) establishing acceptable criteria for datasets to be used as Resources of the NAIRR;

“(C) identifying and providing access to existing curated datasets of value and interest to the NAIRR user community;

“(D) setting up an artificial intelligence data commons to facilitate community sharing and curation of data, code, and models; and

“(E) coordinating as practicable with the National Secure Data Service to make Federal statistical data available to NAIRR users.

“(3) Educational tools and services, including by—

“(A) facilitating and curating educational and training materials; and

“(B) providing technical training and user support.

“(4) AI testbeds, including by—

“(A) facilitating access to artificial intelligence testbeds through which researchers can measure and benchmark engineering or algorithmic developments; and

“(B) developing a comprehensive catalog of open AI testbeds.

“SEC. 5605. NAIRR PROCESSES AND PROCEDURES.

“(a) USER SELECTION.—

“(1) ELIGIBLE USERS.—A researcher, educator, or student based in the United States and affiliated with the following types of entities, if such entity is based in the United States, shall be eligible for access to the NAIRR:

“(A) An institution of higher education.

“(B) A nonprofit institution (as such term is defined in section 4 of the Stevenson-Wydler Technology Innovation Act of 1980 (15 U.S.C. 3703)).

“(C) An Executive agency.

“(D) A federally funded research and development center.

“(E) A small business concern (as such term is defined in section 3 of the Small Business Act (15 U.S.C. 632), notwithstanding section 121.103 of title 13, Code of Federal Regulations⁴) that has received funding from an Executive agency, including through the Small Business Innovation Research Program or the Small Business Technology Transfer Program (as described in section 9 of the Small Business Act (15 U.S.C. 638)).

“(F) A category of entity that the Director of the National Science Foundation and the Director of the Initiative Office, after consultation with the Steering Committee, appropriate Advisory Boards, and the public, determine shall be eligible.

“(G) A consortium composed of entities described in subparagraphs (A) through (F).

“(2) USER ACCESS SELECTION.—The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Steering Committee, shall establish—

“(A) an application for eligible users to request access to the NAIRR; and

“(B) multiple selection processes, with increased scrutiny for an application based on the value or type of Resources of the NAIRR requested.

⁴ The “notwithstanding” provision waives a requirement that often otherwise exempts from the definition of small business concern, as applied by regulation, startups funded by certain private funders (e.g., venture capitalists).

“(b) PRIVACY, ETHICS, CIVIL RIGHTS AND CIVIL LIBERTIES, SAFETY, AND TRUSTWORTHINESS.—The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Ethics Advisory Board, User Committee, Steering Committee, and heads of relevant Executive agencies, shall establish requirements, a review process for applications, and a process for auditing Resources of the NAIRR and research conducted using Resources of the NAIRR on matters related to privacy, ethics, civil rights and civil liberties, safety, and trustworthiness of artificial intelligence systems developed using Resources of the NAIRR. The head of the Program Management Office shall ensure such requirements and process are consistent with policies of relevant Executive agencies.

“(c) SCIENTIFIC INTEGRITY.—

“(1) IN GENERAL.—The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Steering Committee, the Ethics Advisory Board, the Director of the Office of Science and Technology Policy, and the public, shall develop—

“(A) policies for addressing concerns related to matters of scientific integrity, including matters related to the effects or impacts of research and potential research enabled by NAIRR; and

“(B) mechanisms for an employee of the Operating Entity, an employee of the Program Management Office, a member of the Steering Committee or an Advisory Board, a researcher or student affiliated with a NAIRR user, an employee of a NAIRR resource provider, an employee of a NAIRR funding agency, or a member of the public to report violations of the policies established under subparagraph (A), including by confidential and anonymous means;

“(2) CONSISTENCY WITH GOVERNMENT POLICIES ON SCIENTIFIC INTEGRITY.—The policies developed in paragraph (1)(A) shall be published in a publicly accessible location on the website of the NAIRR. Such policies shall, to the degree practicable, be consistent with the Presidential memorandum entitled “Restoring Trust in Government Through Scientific Integrity and Evidence-Based Policymaking,” dated January 27, 2021, or successor document, and reports produced pursuant to such Presidential memorandum (including the report entitled “Protecting the Integrity of Government Science” published by the National Science and Technology Council and dated January 2022, or successor document).

“(3) PUBLIC REPORTING.—The head of the Program Management Office, acting through the Director of the Operating Entity, shall publicly list, and update not less frequently than once every 3 months, the following information about each project receiving any support from NAIRR:

“(A) Project name, description, and anticipated value to the public.

“(B) Names and affiliations of each researcher or student associated with the project.

“(C) A description of data being used for the project.

“(D) Research questions and methods.

“(E) Anticipated reports or other deliverables and associated expected dates for such reports or deliverables.

“(d) SYSTEM SECURITY AND USER ACCESS CONTROLS.— The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Steering Committee, Director of the National Institute of Standards and Technology, and the Director of the Cybersecurity and Infrastructure Security Agency—

“(1) shall establish minimum security requirements for all persons interacting with the NAIRR, consistent with the most recent version of the Cybersecurity Framework, or successor document, maintained by National Institute of Standards and Technology; and

“(2) may establish tiers of security requirements and user access controls beyond the minimum requirements relative to security risks;

“(e) FEE SCHEDULE.—The head of the Program Management Office, acting through the Director of the Operating Entity, may establish a fee schedule for access to NAIRR. The Operating Entity may only charge fees in such fee schedule. Such fee schedule—

“(1) may differ by type of eligible user;

“(2) shall include a free tier of access based on appropriated funds and anticipated costs and demand; and

“(3) may include cost-based charges for—

“(A) persons not otherwise considered eligible users to purchase; and

“(B) eligible users to purchase Resources of the NAIRR beyond those included in a free or subsidized tier;

“(f) OPEN SOURCE AND PUBLIC ACCESS.—The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Science Advisory Board, Technology Advisory Board, the Director of the Office of Science and Technology Policy, and the Director of the Cybersecurity and Infrastructure Security Agency, shall establish policies to encourage—

“(1) principles of open source, including by encouraging software developed for the administration of the NAIRR or using Resources of the NAIRR to be open-source software; and

“(2) to ensure public access of research conducted using Resources of the NAIRR, consistent with the principles outlined in Memorandum on “Ensuring Free, Immediate, and Equitable Access to Federally Funded Research” released by the Office of Science and Technology Policy and dated August 25, 2022, or successor document.

“(g) ENVIRONMENTAL SUSTAINABILITY.—The head of the Program Management Office, acting through the Director of the Operating Entity and in consultation with the Administrator of the Environmental Protection Agency, may establish policies to—

“(1) measure and manage discarded hardware and other electronic waste;

“(2) consider environmental impact of hardware when acquiring, developing, or promoting hardware;

“(3) identify or develop application-development tools that assist NAIRR users in creating energy-efficient applications; and

“(4) research and fund research to study environmental impacts of artificial intelligence systems.

“SEC. 5606. AUTHORIZATION OF APPROPRIATIONS.

“There are authorized to be appropriated to carry out the activities described in this title \$440,000,000 for each of the fiscal years 2023, 2024, 2025, 2026, 2027, and 2028.”

Appendix K. Abbreviations

AI	artificial intelligence
CPU	central processing unit
DEIA	diversity, equity, inclusion, and accessibility
DOE	Department of Energy
DUA	data use agreement
FedRAMP	Federal Risk and Authorization Management Program
FFRDC	federally funded research and development center
FY	Fiscal Year
GIS	geographic information system
GPU	graphics processing unit
HPC	high-performance computing
JSON	JavaScript Object Notation
KPI	key performance indicator
ML	machine learning
NAIIO	National Artificial Intelligence Initiative Office
NAIRR	National Artificial Intelligence Research Resource
NASA	National Aeronautics and Space Administration
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development
NOAA	National Oceanic and Atmospheric Administration
NSDS	National Secure Data Service

NSF	National Science Foundation
OSS	open source software
OSTP	Office of Science and Technology Policy
R&D	research and development
RFI	request for information
RFP	request for proposal
SAP	Standard Application Process
SBIR	Small Business Innovation Research
STEM	science, technology, engineering, and mathematics
STTR	Small Business Technology Transfer

Appendix L. Notes

- ¹ Throughout this report, “AI R&D” is inclusive of foundational AI R&D, use-inspired AI R&D, and translational AI R&D. That is, the NAIRR is relevant not only for researchers advancing the field of AI itself (i.e., foundational research) but also for those who are advancing AI with a use case in mind (i.e., use-inspired research), as well as for those translating AI discoveries and innovations to the market and society (i.e., translational research).
- ² For an alternative, yet compatible, definition of AI, please see the John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115- 232, 132 Stat. 1697, (2018).
- ³ National Artificial Intelligence Initiative Act of 2020 (Pub.L. 116-283) § 5106(a)(1)(A), 15 U.S.C. § 9415(a)(1)(A).
- ⁴ Center for Security and Emerging Technology, “AI Faculty Shortages,” (July 2022), <https://cset.georgetown.edu/wp-content/uploads/CSET-AI-Faculty-Shortages.pdf>.
Additional analysis conducted for the Task Force by the Science and Technology Policy Institute identified 404,858 unique researchers affiliated with U.S. institutions who had published at least one AI-related publication between 2016 and 2021; of these, 14,619 were identified (from any academic department) with at least five AI-related publications.
- ⁵ The percentages listed correspond to the share of computer science, computer engineering, and information PhD recipients in North America whose specialties are known that specialized in “Artificial Intelligence/Machine Learning” or “Robotics/Vision,” as reported in: Stuart Zweben and Betsy Bizot, “2021 Taulbee Survey,” (May 2022), <https://cra.org/wp-content/uploads/2022/05/2021-Taulbee-Survey.pdf>.
- ⁶ Kate Crawford et al., “The AI Now Report: The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term,” (July 2016), https://ainowinstitute.org/AI_Now_2016_Report.pdf.
- ⁷ Stuart Zweben and Betsy Bizot, “2020 Taulbee Survey,” (May 2021), <https://cra.org/wp-content/uploads/2021/05/2020-CRA-Taulbee-Survey.pdf>;
U.S. Census Bureau, “U.S. Census Bureau QuickFacts: United States,” Accessed November 11, 2022, <https://www.census.gov/quickfacts/fact/table/US/RHI725221>.
- ⁸ Daniel Zhang et al., “The AI index 2022 annual report,” (March 2022), <https://aiindex.stanford.edu/report/>.
- ⁹ Ruha Benjamin, “Race after Technology: Abolitionist Tools for the New Jim Code,” Polity, (2019);
Joy Buolamwini and Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, in Conference on Fairness, Accountability and Transparency, pp. 77–91, Proceedings of Machine Learning Research, 2018, <https://proceedings.mlr.press/v81/buolamwini18a.html>;
Kate Crawford, *The Atlas of AI*, Yale University Press, 2021;
Kate Crawford and Trevor Paglen, *Excavating AI: The Politics of Images in Machine Learning Training Sets*, Liverpool Biennial, 9, (2019), <https://www.biennial.com/journal/issue-9/excavating-ai-the-politics-of-images-in-machine-learning-training-sets>;
Catherine D'Ignazio and Lauren F. Klein, “Data Feminism,” MIT press, 2020;
Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, St. Martin's Press, 2018;
Su Lin Blodgett et al., “Language (Technology) is Power: A Critical Survey of 'Bias' in NLP,” *arXiv preprint arXiv:2005.14050*, (2020), <https://arxiv.org/abs/2005.14050>;
Safiya Umoja Noble, *Algorithms of Oppression*, New York University Press, 2018;
Meredith Whittaker et al., “Disability, Bias, and AI,” AI Now Institute, (2019), <https://ainowinstitute.org/disabilitybiasai-2019.pdf>.

-
- ¹⁰ National Security Commission on Artificial Intelligence, “Final Report,” (March 2021), <https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf>.
- ¹¹ Nathan Benaich and Ian Hogan, “State of AI Report 2022,” (October 2022), <https://www.stateof.ai/>.
- ¹² For example, the White House Office of Science and Technology Policy’s October 2022 Blueprint for an AI Bill of Rights defines a path forward for American leadership in the responsible use of AI; see reference 18 and the accompanying announcement of relevant actions across the Federal Government: <https://www.whitehouse.gov/ostp/news-updates/2022/10/04/fact-sheet-biden-harris-administration-announces-key-actions-to-advance-tech-accountability-and-protect-the-rights-of-the-american-public/>.
- ¹³ CHIPS Act of 2022, Pub. L. 117-167, 136 Stat. 1372 (2022).
- ¹⁴ National Artificial Intelligence Research Resource Task Force, “Envisioning a National Artificial Intelligence Research Resource (NAIRR): Preliminary Findings and Recommendations,” (May 2022), <https://www.ai.gov/wp-content/uploads/2022/05/NAIRR-TF-Interim-Report-2022.pdf>.
- ¹⁵ The Networking & Information Technology R&D Program and the National Artificial Intelligence Initiative Office, “Supplement to the President’s FY 2023 Budget,” (November 2022), <https://www.nitrd.gov/pubs/FY2023-NITRD-NAIIO-Supplement.pdf>.
- ¹⁶ As described in: National Research Council, Cooperative Stewardship: Managing the Nation’s Multidisciplinary User Facilities for Research with Synchrotron Radiation, Neutrons, and High Magnetic Fields, The National Academies Press, (1999), <https://doi.org/10.17226/9705>.
- ¹⁷ Note: The Task Force’s interim report proposed that the NAIRR be governed by a Board of Directors. The Task Force has since decided that the NAIRR Operating Entity and resource providers should be primarily responsible to the Steering Committee, and that oversight from separate groups—the Steering Committee and a Board of Directors—could result in conflicting direction to the Operating Entity. Thus, the Operating Entity will be led by its executive staff, answerable to the Steering Committee via the Program Management Office, and with external input from the User Committee, advisory boards, and external evaluator.
- ¹⁸ White House Office of Science and Technology Policy, “Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People.” (October 2022), <https://www.whitehouse.gov/ostp/ai-bill-of-rights>.
- ¹⁹ National Institute of Standards and Technology, “AI Risk Management Framework,” (September 2022), <https://www.nist.gov/itl/ai-risk-management-framework>.
- ²⁰ The White House, “Memorandum on Restoring Trust in Government through Scientific Integrity and Evidence-based Policymaking,” (January 2021), <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/01/27/memorandum-on-restoring-trust-in-government-through-scientific-integrity-and-evidence-based-policymaking/>.
- ²¹ National Science and Technology Council, “A Framework for Federal Scientific Integrity Policy and Practice,” (January 2023), <https://www.whitehouse.gov/wp-content/uploads/2023/01/01-2023-Framework-for-Federal-Scientific-Integrity-Policy-and-Practice.pdf>.
- ²² White House Office of Science and Technology Policy, “Memorandum for the Heads of Executive Departments and Agencies,” (August 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/08/08-2022-OSTP-Public-Access-Memo.pdf>.
- ²³ National Research Platform, “Designed for Growth and Inclusion,” NRP, accessed January 10, 2023, <https://pacificresearchplatform.org/updates/national-research-platform/>.
- ²⁴ Such as FedRAMP, FISMA moderate, or CMMC level 3. See:
General Services Administration, “FedRAMP,” accessed January 10, 2023, <https://www.fedramp.gov/>;
Cybersecurity & Infrastructure Security Agency, “Federal Information Security Modernization Act,” accessed January 10, 2023, <https://www.cisa.gov/federal-information-security-modernization-act>;
Acquisition & Sustainment Office of the Under Secretary of Defense, “Cybersecurity Maturity Model Certification CMMC 2.0,” accessed January 10, 2023, <https://www.acq.osd.mil/cmmc/>.

-
- ²⁵ Department of Health and Human Services Office for Civil Rights, “HIPAA Administrative Simplification,” (March, 2013), <https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/combined/hipaa-simplification-201303.pdf>;
- Defense Counterintelligence and Security Agency, “Controlled Unclassified Information,” accessed January 10, 2023, <https://www.dcsa.mil/mc/isd/cui/>.
- ²⁶ National Science Foundation, “Cyberinfrastructure for Sustained Scientific Innovation (CSSI),” (September 2022), <https://beta.nsf.gov/funding/opportunities/cyberinfrastructure-sustained-scientific-innovation-cssi>.
- ²⁷ National Science Foundation “Pathways Enable Open Source Ecosystems (POSE),” (February 2022), <https://beta.nsf.gov/funding/opportunities/pathways-enable-open-source-ecosystems-pose>.
- ²⁸ Energy Star, “Data Centers,” accessed January 17, 2023, https://www.energystar.gov/products/data_centers
- ²⁹ National Science Foundation, “Standard Application Process,” National Center for Science and Engineering Statistics, <https://nces.nsf.gov/about/standard-application-process>;
- Congress, “H.R.4174 - Foundations for Evidence-Based Policymaking Act of 2018,” (January 2019), <https://www.congress.gov/bill/115th-congress/house-bill/4174>.
- ³⁰ A recent census of the machine learning, artificial intelligence, and data landscape lists more than 500 companies, startups, and tools—many of them open source. Matt Turck, “Red Hot: The 2021 Machine Learning, AI and DATA (MAD) Landscape,” (September 2021), <https://mattturck.com/data2021/>.
- ³¹ For example, the Extreme Scale Software Scientific Software Stack developed by DOE’s Exascale Computing Initiative. E4S, “Home,” accessed January 10, 2023, <https://e4s-project.github.io/>.
- ³² Cnvr.io, “Blueprints,” accessed January 10, 2023, <https://cnvr.io/blueprints/>.
- ³³ National Science and Technology Council, “Desirable Characteristics of Data Repositories for Federally Funded Research,” (May 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/05/05-2022-Desirable-Characteristics-of-Data-Repositories.pdf>.
- ³⁴ ISO, “ISO/IEC 21778:2017 Information technology — The JSON data interchange syntax,” accessed December 21, 2022, <https://www.iso.org/standard/71616.html>.
- ³⁵ Neo4j, “Neo4j Graph Data Platform,” accessed December 21, 2022, <https://neo4j.com/>.
- ³⁶ This is required by the Confidential Information Protection and Statistical Efficiency Act (CIPSEA), Pub. L. No. 107–347, 116 Stat. 2962 (2002).
- ³⁷ Roughly half of the 326,000 entries in data.gov are from State and local governments.
- ³⁸ The Berkeley Data Stack is a collection of open source tools provided to faculty and students across UC Berkeley. University of California, Berkeley, “Berkeley Data Stack,” accessed January 10, 2023, <https://data.berkeley.edu/academics/campus-resources/berkeley-data-stack>.
- ³⁹ National Science Foundation, “Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support,” ACCESS, accessed January 10, 2023, <https://access-ci.org/>.
- ⁴⁰ Such as Kubernetes, SAGE, HTCondor, Kepler, and Pegasus.
- ⁴¹ National Science and Technology Council, “Lessons Learned from Federal Use of Cloud Computing to Support Artificial Intelligence Research and Development,” (July 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/07/07-2022-Lessons-Learned-Cloud-for-AI-July2022.pdf>.
- ⁴² National Institutes of Health, “SRA in the Cloud,” National Library of Medicine, accessed January 10, 2023, <https://www.ncbi.nlm.nih.gov/sra/docs/sra-cloud/>.
- ⁴³ National Aeronautics and Space Administration, “Earth Data Cloud Evolution,” Earth Data, (March 2022), <https://earthdata.nasa.gov/eosdis/cloud-evolution>.
- ⁴⁴ U.S. Digital Service, “Digital Services Playbook,” accessed January 10, 2023, <https://playbook.cio.gov/>.
- ⁴⁵ VanEseltine et al. “Report on the Scientific and Economic Effects of XSEDE” (August, 2020), Institute for Research on Innovation & Science (IRIS), University of Michigan, https://iris.isr.umich.edu/wp-content/uploads/2022/12/IRIS_XSEDE_pilot_report_12August2020_for_comment.pdf.

-
- ⁴⁶ National Institutes of Health, “STRIDES Initiative,” Office of Data Science Strategy, accessed January 10, 2023, <https://datascience.nih.gov/strides>.
- ⁴⁷ Cloudbank, “Managed Services to Simplify Cloud Access for Computer Science Research and Education,” accessed January 10, 2023, <https://www.cloudbank.org/>.
- ⁴⁸ National Science Foundation, “PATH,” accessed January 10, 2023, <https://path-cc.io/>.
- ⁴⁹ OpenAI, “OpenAI API,” accessed January 10, 2023, <https://openai.com/api/>.
- ⁵⁰ National Institutes of Health, “All of Us Research Program,” <https://allofus.nih.gov/>.
- ⁵¹ National Institutes of Health, “National COVID Cohort Collaborative,” accessed January 10, 2023, <https://ncats.nih.gov/n3c>.
- ⁵² These are the shorthand names of major international AI conferences: the annual conference of the Association for the Advancement of Artificial Intelligence (AAAI), Conference on Neural Information Processing Systems (NeurIPS), the Association for Computing Machinery Special Interest Group on Knowledge Discovery and Data Mining Conference (KDD), The Association for Computing Machinery and IEEE Computer Society’s Supercomputing conference (SC), and the International Conference on Machine Learning (ICML).
- ⁵³ International Science Council Committee on Data, “Home- CODATA, The Committee on Data for Science and Technology,” accessed January 10, 2023, <https://codata.org/>.
- ⁵⁴ These figures are in 2022 dollars and do not account for inflation or decline in graphics processing unit (GPU) prices over time.
- ⁵⁵ Or Sharir et al., “The Cost of Training NLP Models,” *arXiv preprint*, arXiv:2004.08900v1, (2020), <https://arxiv.org/pdf/2004.08900.pdf>.