Agenda

11:00-11:10	**Welcome and Administrative Remarks,** Lynne Parker & Manish Parashar
11:10-12:10	**Readout and Discussion of Draft Recommendations: User Access Controls and Usable Security,** Elham Tabassi
12:10-1:10	**Readout and Discussion of Draft Recommendations: Privacy, Civil Rights, & Civil Liberties,** Manish Parashar
1:10-1:40	**Break**
1:40-2:40	**Readout and Discussion of Draft Recommendations: Technical Integration,** Mike Norman
2:40-2:55	**Briefing: Public-Private Partnerships and Considerations for the NAIRR,** Emily Grumbling, Lisa Van Pay, & Morgan Livingston, STPI
2:55-3:55	**Panel: User Perspectives on the NAIRR** • Tom Dietterich, *Distinguished Professor Emeritus in the Collaborative Robotics and Intelligent Systems Institute, 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 of Electrical and Computer Engineering, Drexel University* • Rima Seiilova-Olson, *Co-Founder and Chief Machine Learning Scientist, Kintsugi* • Carlos Theran-Suarez, *Instructor, Computer and Information Sciences Department, Florida A&M University*
3:55-4:10	**Discussion: Public-Private Partnerships and Sustainment Considerations for the NAIRR,** Lynne Parker
4:10-4:25	**Break**
4:25-5:00	**Discussion: Defining Indicators of Success for the NAIRR,** Manish Parashar
5:00-5:45	**Discussion: Interim Report Outline and Next Steps in Recommendation Finalization,** Lynne Parker
5:45-5:55	**Questions from Public,** Manish Parashar
5:55-6:00	**Closing Remarks,** Manish Parashar

Usable Security and User Access Controls WG

Draft Recommendations

February 2022

Elham Tabassi, Julia Lane, Andrew Moore

Charge to Working Group

- What security requirements should be built into the technical architecture and governance of the NAIRR? What do acceptable protections look like?
- What user access controls should the NAIRR implement? Should these controls vary in accordance with the resource being accessed by the user, and if yes, how so? How about by the type of user?
- What kinds of training resources should be made available to users to educate them on the security and access-control policies and processes?
- How should security and access controls be monitored and updated by a NAIRR management entity?

Process

- Met weekly via Zoom
- Developed 2-page overview of preliminary recommendations
- Considered the Administrative Data Research Facility model
- Some open questions remain for TF discussion

Findings

Finding 1: The NAIRR can leverage the results of extensive efforts in the public and private sector to develop security processes and policies for cyberinfrastructure, e.g.:

- NIST and ISO standards
- Sector-specific frameworks such as Trusted CI
- Federal Risk Authorization and Management Program (FedRAMP)
  - A government-wide program
  - Promotes the adoption of secure cloud services by providing a common security framework
  - Familiar to government entities and contractors

Finding 2: There is heterogeneity in the security needs for different NAIRR assets; some information assets require stronger protections than others. For example:

- Data collected on human subjects is often confidential
- Much government data is protected by statute
- Research could involve confidential intellectual property (IP)
- Compute environments should be protected from unauthorized uses, including by bad actors

Finding 3: Cybersecurity is a rapidly changing field; security controls, mechanisms, and policies will need to keep up with new developments

- New threats evolve constantly, posing new risks
- Without expert, trained, and sufficiently-compensated staff, NAIRR security objectives will not be met

Finding 4: NAIRR security risks are as much human as technical

- User behaviors affect system security
- NAIRR policies are as important as technical controls

Assumptions

- Data and compute resources for a given project will reside on the same platform
- The NAIRR should accommodate both open science and research activities that warrant restricted access (e.g., with PII or other sensitive data)
- The NAIRR will provide access through a single sign-on

Recommendations

Recommendation 1: Leverage FedRAMP standards (not approval process) initially, transitioning to NAIRR-specific standards over 2 years.

- The NAIRR will need to make substantial investments in establishing technical controls as well as governance to determine the policies for how data arrive on a platform, and how they are allowed to leave.
- The NAIRR shall assemble an advisory group to develop a NAIRR-dedicated FedRAMP-looking solution and plan for its regular monitor and update. The advisory group should have expertise/background in the following: Security, operational consideration, risk management, scientific research consideration (security vs usability)
- Focus on incentive structure rather than enforcement structure
- Include regular monitoring and update of the security controls as part of NAIRR governance.

Recommendation 2: Implement a tiered access model.

- The heterogeneity in protection requirements will require different levels of permissions for individuals to work at different levels of access.
- This approach is consistent with the Foundations for Evidence-based Policymaking Act (Public Law No: 115-435)
- FedRAMP standards act as the floor, not the ceiling. Full FedRAMP approval may be required for certain confidential information assets. Elements with lower risk levels need not obtain full FedRAMP approval and could adopt only the needed aspects of FedRAMP standards.
- The high bar of FedRAMP approval could be prohibitive to potential resource providers. A flexible, tiered approach can help open participation to different resource providers, while supporting high security where needed.
- The NAIRR should provide access to a “green-zone” to all users, the lowest tier of access; other users will have access to secure enclaves for working with highly sensitive data or algorithms.

Recommendation 3: Provide regular and continuous hands-on training for NAIRR staff and users

- To reduce human error, users should receive some basic training on system security policies and processes, including:
  a. NAIRR system protections– *this will help to promote trust in the system*
  b. User roles and responsibilities– *for awareness of risks and consequences*
  c. Step-by-step instructions to guide users through key actions
  d. Different tiers of training appropriate to different tiers of access
- The initial training could be offered using multiple media– e.g., videos, step-through computer-based training, PDFs – and accessible to individuals with differing abilities.
- Users should also have access to 24-7 live security support

Recommendation 4: The NAIRR System should be regularly monitored, and its security controls periodically updated as part of its management and governance.

Recommendation 5: The NAIRR System should have dedicated, expert technical/security staff

- Staff should be compensated competitively
- Staffing needs may have implications for, or depend on, the ownership, administration, and governance model for the NAIRR

Open Questions

- Are there any scientific resource/applications/services that would not be supported by requiring FedRAMP approval or compliance with FedRAMP standards?
- Can there be a model where the NAIRR portal provides access to both FedRAMP-approved and non-approved services, or would there need to be separate portals?
- What entity will sign off on NAIRR Authorization to Operate (ATO)?
- What approach is needed to endpoint protection? (User system requirements?)
- How will NAIRR balance security needs of logging user activities with real and perceived concerns regarding academic freedom?

Privacy, Civil Rights, and Civil Liberties

Working Group Members:

Daniela Braga, Oren Etzioni, Manish Parashar, Fred Streitz

Framing Questions

- What governance policies should the NAIRR implement for projects that use NAIRR resources to promote fair, trustworthy, and responsible AI research and development? For example, what constraints, if any, should be placed on NAIRR use to guard against unethical uses of the NAIRR? To ensure the responsible conduct of research performed via the NAIRR?
- How should bias, fairness, and privacy considerations be addressed through NAIRR data-governance policies? For example, what constraints, if any, should be placed on NAIRR data resources and use to protect individual privacy?
- How should the NAIRR governance and management approaches align with policies established by Federal funding agencies, as related to privacy, civil rights, and civil liberties?
- How should NAIRR resource allocation and usage policies address civil rights and civil liberties considerations?
- What educational and training resources should be provided through the NAIRR to reinforce privacy, civil rights, and civil liberties?

Understanding issues, requirements, and potential solutions for the NAIRR related to privacy, civil rights and civil liberties

- Conversations with experts (Wednesday, 01/26/2022)
  - Prof. Kate Crawford, Founding Director, AINow (NYU)
  - Prof. Aaminah Norris, Director of Education, Algorithmic Justice League
  - Dr. Harlan Yu, Executive Director, Upturn
- Discussion and recommendations from the panel at the 12/13/21 NAIRR TF meeting

Key highlights from discussion on 01/26/2022

- NAIRR has an opportunity to create strong governance frameworks, standards for data governance and stewardship
- There is a need for a repository of datasets and their status and provenance, audits of datasets in active use, as well as mechanisms for deprecating datasets (and applications) as appropriate
- The evaluation of AI research/tools cannot be done in isolation and needs to consider the broader social, political, and historical context
- Transparency (and communication) provides an important baseline to start thinking about how to mitigate harms

NAIRR strategy for protecting privacy, civil rights and civil liberties

- **Overarching strategy**: Transparency and Oversight
- **Vetting:** Creating a framework for vetting what research gets done on NAIRR, including developing user agreements, user/resource-provider policies, and proposal/resource evaluation processes
- **Operations:** Mechanisms to support security and privacy requirements, enable monitoring, and ensure compliance with governance policies
- **Reporting:** Regular reporting requirements for stakeholders
- **Training:** Training to make users and resource providers aware of policies, expectations, responsibilities, and best practices

"civil liberties" = protected by Constitutional Bill of Rights; "civil rights" = protected by law

Vetting research enabled by the NAIRR

- Assumption: Projects running on the NAIRR will focus on AI-related research and not production AI services
- Integrate privacy, civil rights, and civil liberties policies and requirement into user agreements governing access to all NAIRR resources
- Integrate privacy, civil rights, and civil liberties policies and requirements into the proposal evaluation process required for accessing the NAIRR
  - Leverage the Ethics Review Board proposed by the Governance Working Group (upcoming slide)
    - The review board should be diverse (by gender, ethnicity, geography, organization types and sizes, etc.) and include relevant expertise and representation for relevant stakeholder communities
  - Leverage existing mechanisms at funding agencies such as NSF, DOE, etc., to ensure research proposals are appropriately vetted
- Ensure transparency of the selection/vetting process

Vetting outcomes of research enabled by the NAIRR

- Establish appropriate reporting requirement
  - Regular project reporting requirements at a to be defined cadence, with explicit focus on privacy, civil rights and civil liberties
  - Regular NAIRR reporting requirements at a to be defined cadence, with explicit focus on privacy, civil rights and civil liberties
- Leverage the Ethics Review Board proposed by the Governance Working Group to evaluate reports and the implication
- Outcomes of the project and methods used should be made public and shared broadly
- Establish mechanisms for receiving and acting on ongoing feedback on results of NAIRR-supported research
  - Continuous monitoring of the use of NAIRR resources and supported research
  - Address longer-term impact of NAIRR resources (e.g., datasets) and research
- Balance reporting overheads with the need for research agility and progress

Governance Working Group Possible Recommendation 2a: Governance Structure

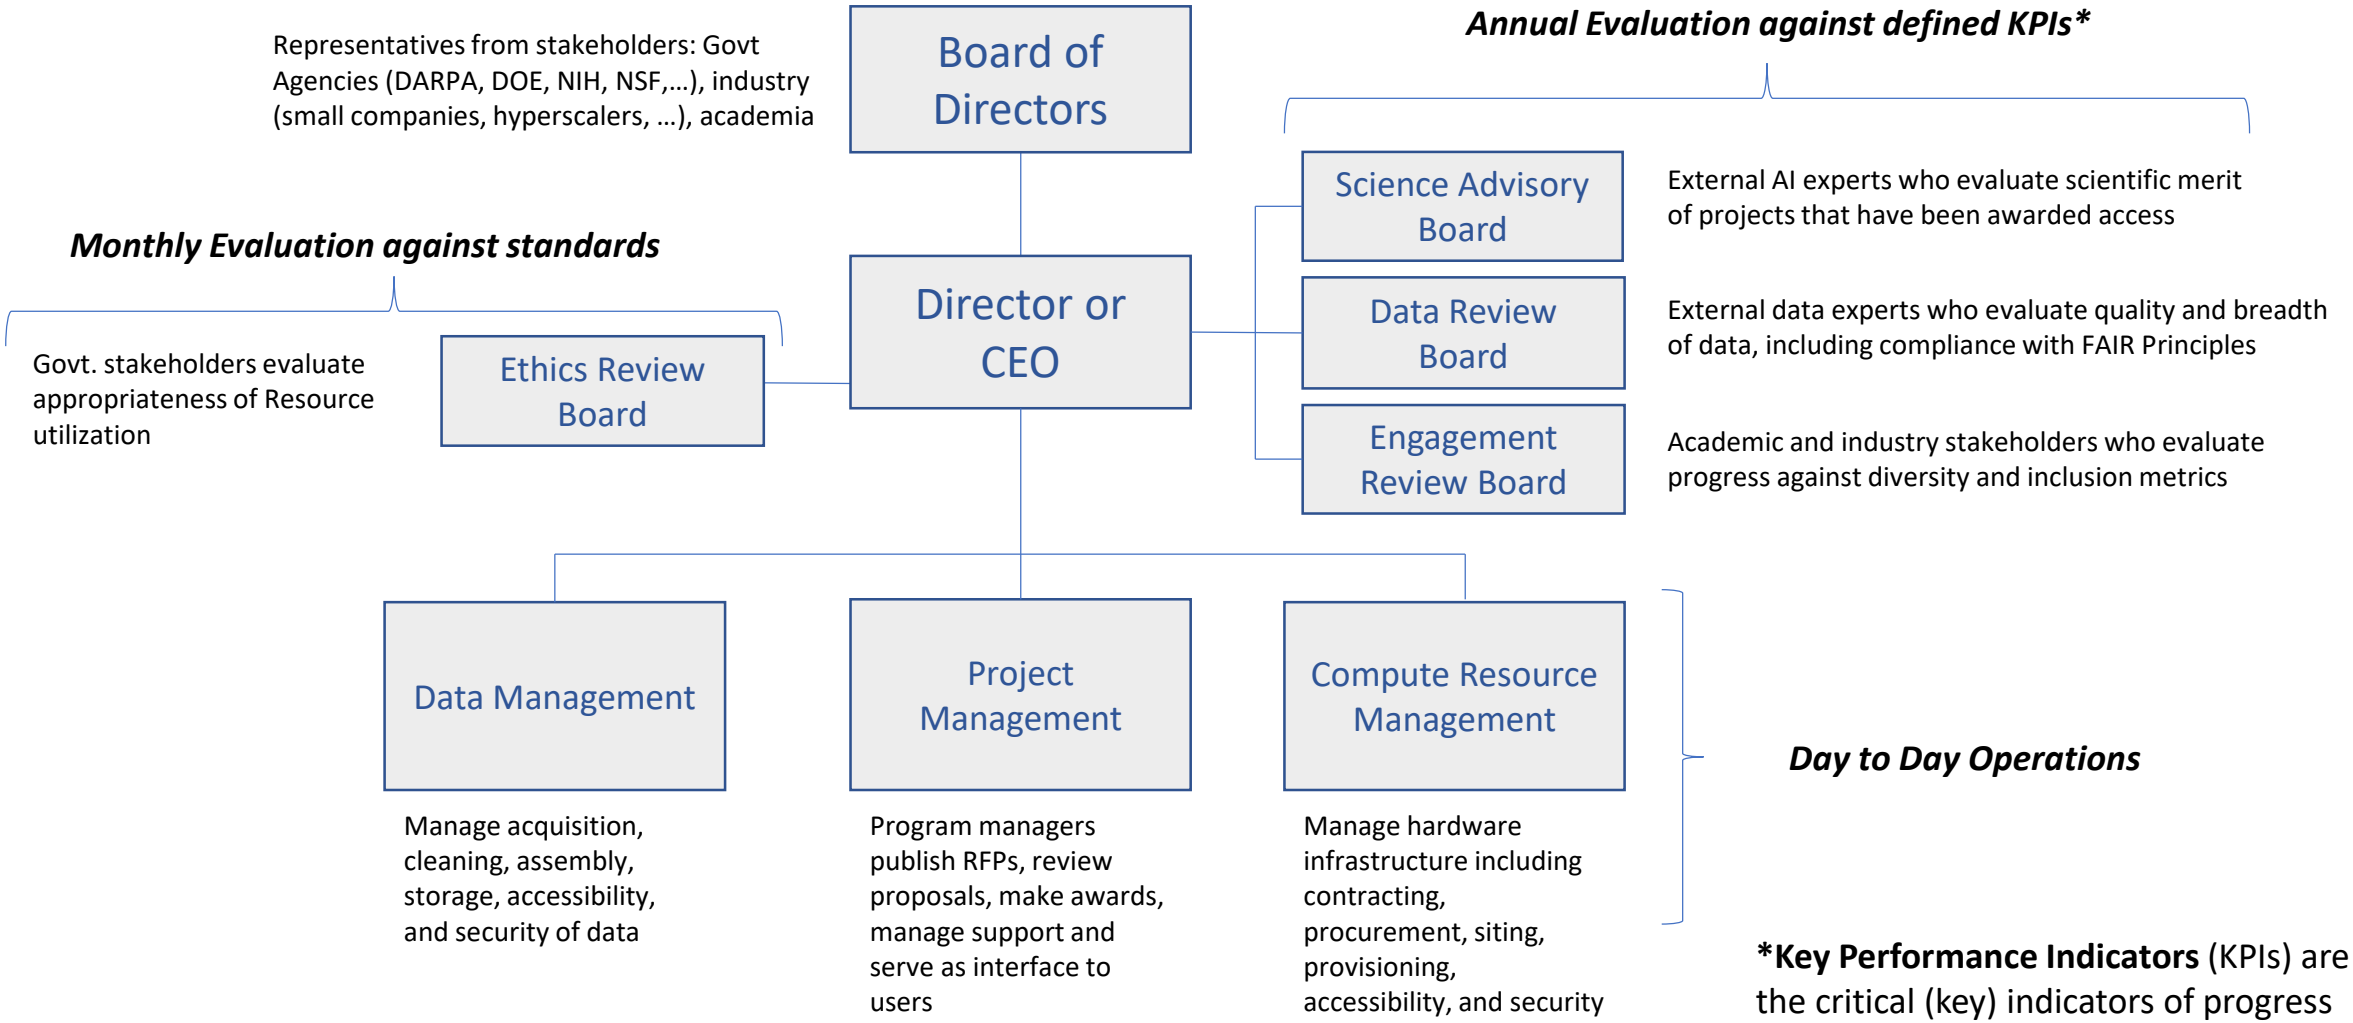

***Key Performance Indicators** (KPIs) are the critical (key) indicators of progress toward an intended result.

Monitoring NAIRR Operations

- Establish acceptance criteria and recommended best practices for all resources joining the NAIRR
- In case of *third-party datasets* made available via the NAIRR, these include
  - Certification standards
  - Trusted/validated testing refence datasets (e.g., an audit system)
  - An inventory of datasets provided by NAIRR and their history/provenance, including a list of deprecated datasets
  - Best practices for datasets (e.g., DOIs) and other resources
- Maintain appropriate security/privacy requirements and mechanisms
  - Work with the Security Working Group and the Data Resources Work Group
- Monitor NAIRR operations; ensure transparency of operations and use
  - Examples of what is monitored: logs of data usage, systems levels, security compliance; application provenance; user logs
  - Monitoring done by NAIRR in coordination with resource providers
  - Support longer-term (e.g., community-based) monitoring and reporting
- Establish protocols, communication infrastructure, and an engagement plan for use in the event that a violation is identified

Training

- Approach to training:
  - All NAIRR users will be required to take compulsory training about rights, responsibilities, and best practices related to privacy, civil rights and civil liberties before being granted access to the NAIRR
  - Training is required to be refreshed annually
  - Content delivered online using e-learning mechanisms
- Developing training content:
  - Content should be developed by third parties with specialization in sensitivity training and in civil rights and civil liberties and how do they apply to AI research, in conjunction with experts in bias and responsible AI

Example Training Resources

- University resources:
  - UC Berkeley's ML Fairness Mini-Bootcamp (includes interactive Jupyter notebooks)
  - Harvard University's Embedded EthiCS modules
  - Stanford's Embedded EthiCS program and publicly available lectures
  - MIT Open Courseware's Social and Ethical Responsibilities of Computing
  - Princeton's AI Ethics Case Studies
  - Some universities also offer courses specific to Natural Language Processing ethics, including Stanford University, Carnegie Mellon University, University of Washington, and UCLA
- Non-university courses:
  - AI Ethics Course
  - Fast AI Practical Data Ethics Course
  - Kaggle Intro to AI Ethics Course
  - Google ML Fairness module
  - AI Bias for K-12 learners

References

- Responsible Computing Research Ethics and Governance of Computing Research and its Applications | National Academies
- fulltext.pdf (rti.org)
- Algorithmic Justice League https://www.ajl.org/
- Gender Shades http://gendershades.org/overview.html
- Voicing Erasure, Algorithmic Justice League https://www.ajl.org/voicing-erasure
- Excavating AI https://excavating.ai/
- Algorithms of Oppression

Technical Integration Working Group Outbrief:

Mike Norman (lead)
Dan Stanzione
Mark Dean
Fei-Fei Li

Charge Questions

1. How should the NAIRR cyberinfrastructure be designed to integrate a **diverse and distributed set of resources** that include computation, data, software workflows, high-speed networking, and AI testbeds?
2. How can the NAIRR be designed so that it can **rapidly adapt** to new requirements, technologies, and technical approaches as they emerge?
3. How should the technical integration of the NAIRR resources be tuned to allow for a **seamless and intuitive user experience** across a spectrum of users? What associated training and documentation is needed?
4. What novel approaches might be needed to effectively **integrate data repositories and edge computing devices**?
5. How should **usage and allocations** be tracked and reported across the NAIRR cyberinfrastructure?
6. What kinds of **staffing and expertise** would be required to set up and maintain the NAIRR cyberinfrastructure?

Design assumptions

- NAIRR will be part of an **integrated national CI ecosystem** that is constantly evolving
- There may be one or more **NAIRR-customized resources** (compute, data) in the ecosystem (production, experimental)
- NAIRR will **leverage ongoing and planned federal agency investments** in integrated CI and interoperate with them
- Other agencies/organizations may wish to contribute **specific/unique resources** to NAIRR (e.g., datasets) which will need to be integrated
- Integration of **AI/ML data repositories, edge computing resources, and AI testbeds** represent new elements to the CI ecosystem, and require special attention for NAIRR to be successful
- NAIRR users should be able to select their AI application, compute resource, and data source(s) from a list and launch and monitor jobs from a **portal providing a uniform, integrated view with a minimum of effort**

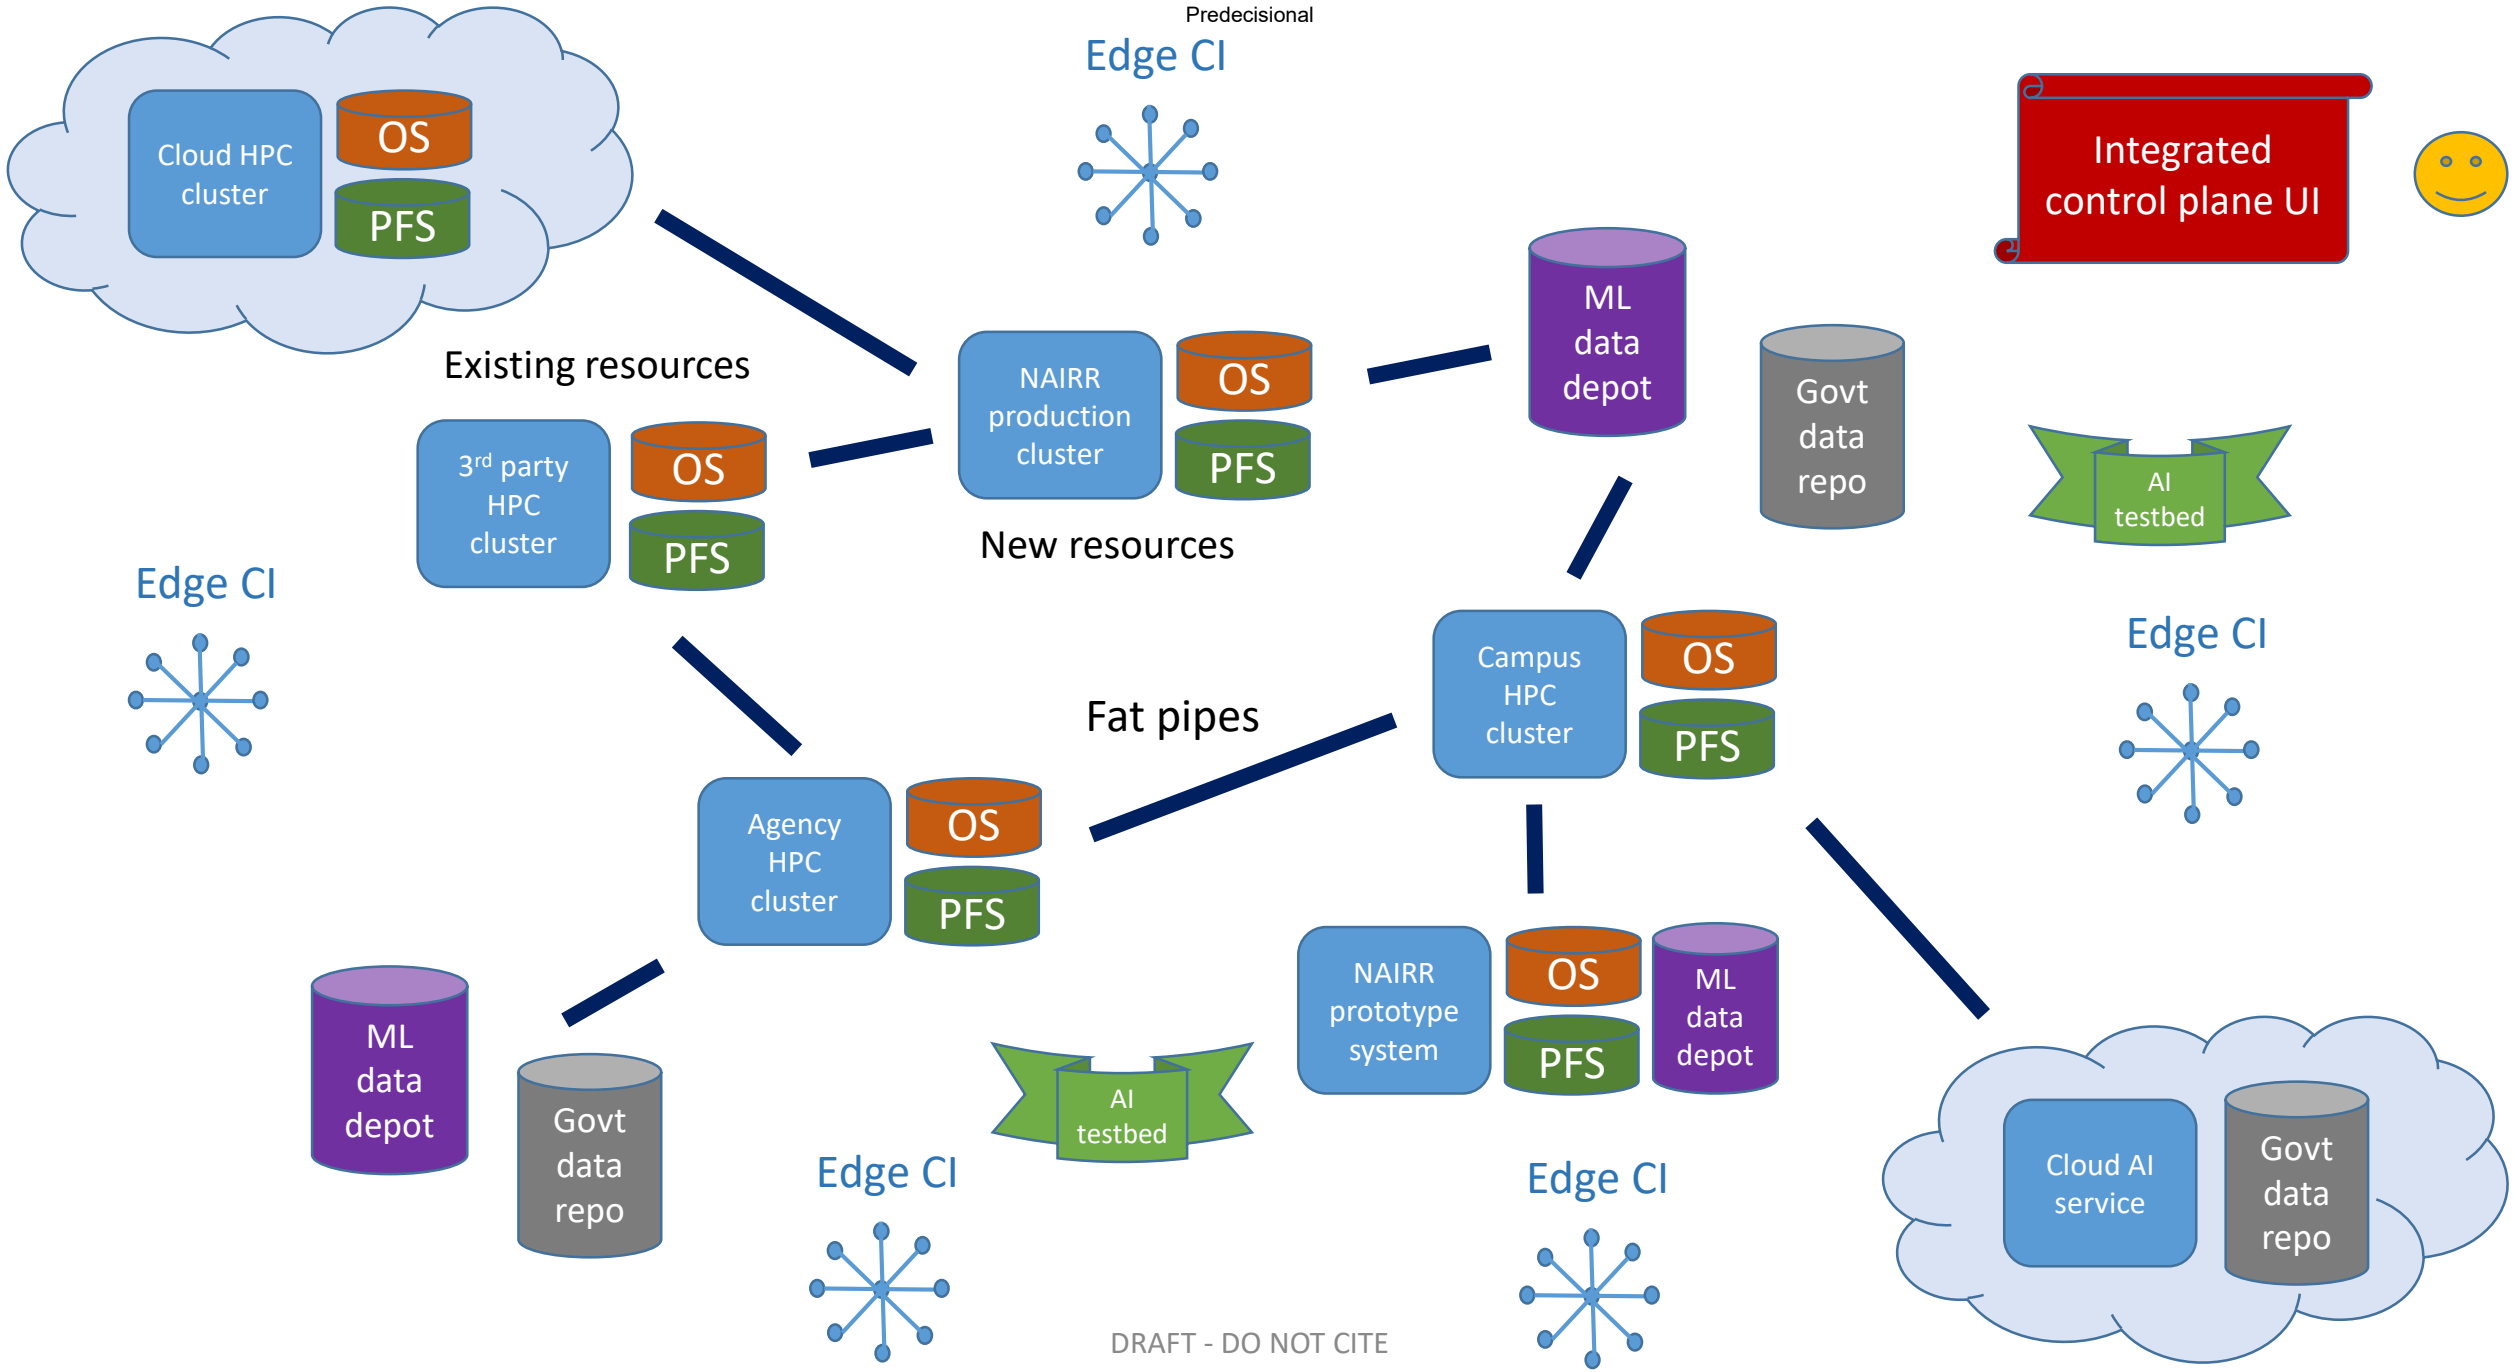

Information Gathering

- In addition to WG members' expertise in scientific computing technology and management, we sought the testimony of

Ian Ferreira, Core Scientific Inc.

technical integration strategies of Plexus™ user portal unifying access to private, multi-cloud, and hybrid cloud resources

Pete Beckman, Argonne Nat'l Lab & U. Chicago

Status update on SAGE edge computing CI project

Key Findings

- Technical integration of federated compute resources is **mature**, providing strategies and solutions for NAIRR to emulate (e.g., XSEDE, OSG)
- Technical integration of AI/ML data repositories, edge computing resources, and AI testbeds is **less mature** and represent new elements to the CI ecosystem requiring special attention for NAIRR to be successful
- NAIRR has the **opportunity to "move the needle"** toward a more data-centric CI ecosystem by supporting **broad adoption of best practices** in scientific data management, discovery, access, and curation. Edge computing and AI testbeds are important drivers for this **transformation**.

1. How should the NAIRR cyberinfrastructure be designed to integrate a diverse and distributed set of resources that include computation, data, software workflows, high-speed networking, and AI testbeds?

- Recommendation

  NAIRR should embrace standards, de-facto standards, and best-of-breed open source solutions wherever possible to ensure a vibrant, growing AI ecosystem. NAIRR should avoid "one-off" integrations.

- Elaboration
  - Integrate federated compute/data/edge/testbed resources behind a single user-facing portal with multiple pathways/interfaces to resources behind it.
  - Resource providers must adhere to standard APIs and/or best practice methodologies where they exist.
  - Support all compute modalities: interactive, batch, on-demand, always-on, and composable systems
  - Co-locate NAIRR ML data depots with compute resources.
  - Support user-supplied apps and data sets with an approval process.
  - User portal designed for distributed, collaborative teams and improved user experience.

1. How should the NAIRR cyberinfrastructure be designed to integrate a diverse and distributed set of resources that include computation, data, software workflows, high-speed networking, and AI testbeds?

- Options to consider
  - Contract out the development of a “single control plane” user portal to external bidder.
  - Adopt Kubernetes as universal resource manager.
  - Incentivize creation of exemplar ML data depots supporting FAIR interfaces and practices.

2. How can the NAIRR be designed so that it can rapidly adapt to new requirements, technologies, and technical approaches as they emerge?

- Recommendation

  NAIRR resource pool should be refreshed frequently based on user need and trends and technological advances.

- Elaboration
  - Modular design, agile development philosophy.
  - Discoverable resources supporting standardized protocols/APIs .
  - Continuously updated resource catalog with search capability.
  - Integrate cloud services and open source technologies where appropriate.
  - Require and enforce resource provider SLAs which may be tiered according to resource maturity.

2. How can the NAIRR be designed so that it can rapidly adapt to new requirements, technologies, and technical approaches as they emerge?

- Options to consider
  - Nurture the development of standardized edge computing middleware and APIs which is currently immature.
  - Define minimum acceptable SLA for AI Testbed RPs

3. How should the technical integration of the NAIRR resources be tuned to allow for a seamless and intuitive user experience across a spectrum of users? What associated training and documentation is needed?

- Recommendation
  - **Design NAIRR user portal as “walk-up tooling” for scientists with consistent user experience across private, multi- and hybrid cloud infrastructures. Support alternate access methods (e.g., shell, scripting) for more advanced users.**

- Elaboration
  - Apps, data sets, and compute resources are high level abstractions in resource catalogs. User portal manages workflow orchestration and monitoring.
  - In-portal help facility.
  - Instrument user portal to track most visited pages and tutorials.
  - Tiered user-training documentation and interactive tutorials (beginner, intermediate, advanced) created by resource provider.
  - Consolidated, searchable training materials catalog (centralized or federated).
  - Portal usability analysis.

3. How should the technical integration of the NAIRR resources be tuned to allow for a seamless and intuitive user experience across a spectrum of users? What associated training and documentation is needed?

- Options to consider
  - Recommender/configuration service to guide selection of resources.
  - Catalog of exemplar use cases (CloudBank).
  - User self-help facilities via NAIRR user forum/Slack/chat channels.

4. What novel approaches might be needed to effectively integrate data repositories and edge computing devices?

- Recommendation
  - **Establish a network of exemplar ML data repositories with powerful search and retrieval capabilities. Encourage the development of standard edge computing middleware**.

- Elaboration
  - Avoid “one-off” integrations.
  - Embrace standards wherever possible.
  - Encourage emergence of standard interfaces to edge and data resources by partnering with best practice efforts.
  - Federated data search and access services with API support.
  - Role-based access to restricted data integrated with search and access service.

4. What novel approaches might be needed to effectively integrate data repositories and edge computing devices?

- Options to consider
  - Survey current best practices in building ML data repositories (e.g., EarthCube Geocodes)
  - Acquire and serve up scientifically and societally important data sets (social, environmental) to encourage NAIRR usage and AI innovation.
  - NAIRR storage caches for big data
  - Support migration of data to publication in data journals (data offramps).

5. How should usage and allocations be tracked and reported across the NAIRR cyberinfrastructure?

- Recommendation
  - **Allocate resources in USD ($) or equivalents.**

- Recommendation
  - **Instrument resources and user portal to automatically track and report both resource utilization and user data.**

- Elaboration
  - CloudBank adopted $ as simplest approach to deal with public cloud pricing and credits
  - HPC resources can be similarly denominated (XSEDE calculates $-equivalent for an allocation)
  - CoreScientific and Rackspace price hybrid cloud resources in $
  - Public cloud usage monitored in near real-time similar to CloudBank with automated overspend alerts.

6. What kinds of staffing and expertise would be required to set up and maintain the NAIRR cyberinfrastructure?

- NAIRR resource providers
  - AI/ML devops teams
  - HPC & storage systems administrators
  - High performance WAN and LAN network specialists
  - Scientific data management/curation experts
  - Site-specific training specialists

- NAIRR user support center
  - AI training specialists & solution architects
  - User portal developers
  - data analysts reporting on NAIRR usage patterns and trends.

6. What kinds of staffing and expertise would be required to set up and maintain the NAIRR cyberinfrastructure?

- Options to consider
  - Outsource construction, operation, and maintenance of NAIRR user portal to external awardee.

Key Findings

- Technical integration of federated compute resources is **mature**, providing strategies and solutions for NAIRR to emulate (e.g., XSEDE, OSG)
- Technical integration of AI/ML data repositories, edge computing resources, and AI testbeds is **less mature** and represent a new element to the CI ecosystem requiring special attention for NAIRR to be successful
- NAIRR has the **opportunity to "move the needle"** toward a more data-centric CI ecosystem by supporting broad adoption of best practices in scientific data management, discovery, access, and curation. Edge computing and AI testbeds are important drivers for this **transformation**.

IDA SCIENCE AND TECHNOLOGY POLICY INSTITUTE

Considerations for Partnerships to Realize the Vision for the NAIRR

Lisa Van Pay
Emily Grumbling
Morgan Livingston

February 16, 2022

Science and Technology Policy Institute
1701 Pennsylvania Ave, NW ● Suite 500 ● Washington, DC 20006-5825

Outline and objectives

- Resources and options for partnerships with NAIRR
- Advantages and limitations of different partnerships
- Areas where partnerships may provide unique value

Objective: Provide an overview of areas where partnerships may help realize the vision for NAIRR, and outline potential advantages or limitations.

We considered resource partnerships in key areas of interest for the NAIRR

Compute

Data

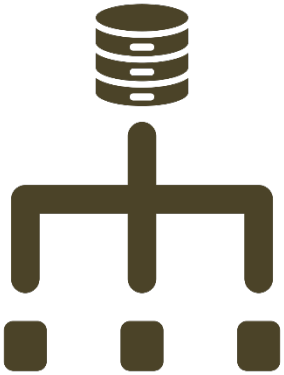

Data Infrastructure

Education/Training Tools

Stakeholder Interests

NAIRR user needs and governance will influence the "fit" of different potential partners

For each resource or component area, the Task Force will need to consider:

Authority of the NAIRR

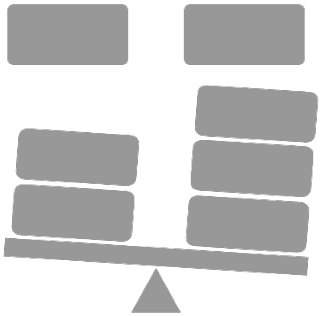

Risks, Responsibilities, and Returns

Capabilities and Scale

Partner Incentives

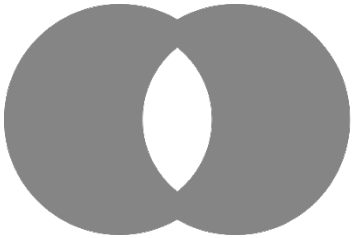

Areas Where Partnerships Can Provide Unique Value

Stakeholder Support

Compute

Option	Advantages	Limitations
NAIRR builds	More control over resource architectures and policies Potentially cost-competitive in the longer term	Slower to launch More costly in the near-term Slower to reflect technology advances Must build large, highly expert staff NAIRR responsible for security and troubleshooting
NAIRR partners with federally funded CI resources	Lower costs in the near term Leadership-class computing resources available Shorter launch time Established R&D infrastructure	Slower to reflect technology advances Resource expansion may be needed to accommodate NAIRR + existing demand
NAIRR partners with commercial CSPs	Shorter launch time Quicker to reflect technology advances Lower costs in the near term Provides users experience on commercial resources relevant to future careers	Constrained by resource types and price points Reinforces reliance on commercial/for-profit resources IP and end user policies may need to be negotiated in detail

Data

Option	Advantages	Limitations
Federal data	Reliable data on pressing social needs	Compliance requirements, lack of resources and mechanisms to share
Private sector data	High volumes of data suited for AI research (e.g., video)	Lack of incentives to share
University and Non-Profit	Datasets specific to research needs	Lack of resources, mechanisms, and clear protections to share

Data Infrastructure

Option	Advantages	Limitations
Built by NAIRR	Trusted data management, customizable	Slow to launch
Leverage open source tools	Cost savings, support open source ecosystem	Vulnerability and quality management
Provided by partner entity	Fast to launch	Greater potential for mistrust, vendor lock in

Educational and Training Tools

Option	Advantages	Limitations
Partner with higher education (Traditional)	Straightforward approach; provides platform for existing educational resources	Unlikely to address lack of diversity across the AI workforce
Partner with training program(s) (Non-Traditional)	May help diversify the AI workforce via recruitment of non-traditional students	May not be sufficient to grow the technical talent pool or develop cutting-edge AI applications or use-cases
Partner with higher ed and training programs (Mix of traditional and non-traditional)	Can offer opportunities that span the workforce pipeline, from skill building and upskilling to higher education	Integrating resources to provide a seamless experience across educational offerings may be difficult

Stakeholder Interests

Option	Advantages	Limitations
Representatives on advisory boards	Subject matter experts in key areas can provide crucial information to inform NAIRR operations and governance	Engaging select individuals may not fully represent the interests of a broader stakeholder group
Include civil society groups in partnerships	Provides a formal way to engage stakeholder groups in partnerships	Successfully balancing the interests of all stakeholders in a given partnership may prove challenging
Develop specific partnerships to address stakeholder issues	Partnerships could specifically address privacy, safety, or oversight; Partnerships would demonstrate to the public that the NAIRR is taking meaningful action to address concerns	Actions or recommendations would need to feed into NAIRR governance or policies to “have teeth”

Partnerships could enable access to key capabilities or experiences for NAIRR users

- Existing cyberinfrastructure offers research and training options not readily replicable in-house
- Collaboration with a range of data holders provides unique opportunities for AI-driven discovery
- Partnerships with top content providers could integrate educational resources across the workforce pipeline
- Integrating partner resources via the NAIRR portal magnifies the environment's utility
- Engaging stakeholders on privacy, safety, or oversight needs

Defining Indicators of Success for the NAIRR

MANISH PARASHAR, OFFICE DIRECTOR, OFFICE OF ADVANCED CYBERINFRASTRUCTURE, NATIONAL SCIENCE FOUNDATION.

Proposed Recommendations

- Indicators of success for a NAIRR should be grounded in progress towards the intended outcomes of establishing and operating the NAIRR:

Innovation	Spur innovative and novel methods in foundational and use-inspired AI research.
Diversity	Lower the barriers so that anyone within the U.S. can contribute to the AI innovation ecosystem.
Capacity	Promote AI skills and knowledge through expanded access to state of the art AI resources.
Ethics	Support the development and adoption of trustworthy and responsible AI

Proposed Recommendations

- The management of the NAIRR should capture and record data related to the investment of resources, activities/usage, outputs, and impacts of the NAIRR.
- These indicators should be regularly assessed to gauge progress towards intended outcomes and inform the management's resource investment and service management decisions, driving a nimble approach that is responsive to the user community.
- The NAIRR management entity should budget sufficient funds for robust data collection and evaluation activities, and design NAIRR infrastructure processes to capture key analytics measures from their outset.
- Efforts to collect the baseline data around measures of composition, strength, and production of the AI research community should begin as soon as feasible, in parallel with the development and launch of the NAIRR.

Proposed Recommendations

- The NAIRR should be externally evaluated at regular intervals to determine the value of the research that is supported by the NAIRR and disseminate this information about the value to the taxpayer, Congress, and the supporting Federal agencies, using open, reproducible, and transparent methods to do so.

- The external evaluator, in conjunction with NAIRR management and the NAIRR advisory board, should identify an appropriate counterfactual for each set of programmatic investments once the scales of these investment are determined in order to gauge the value of the NAIRR's role in advancing the AI research ecosystem.

Proposed Recommendations

- The NAIRR management organization should capture four levels of performance/success indicators: 1) measures associated with resource investments; 2) measures of resource usage/activities; 3) measures of outputs; and 4) measures of impact.

- These measures should be complemented by annual user surveys and ongoing user engagement mechanisms to gauge user satisfaction, assess usability of the platform, identify emerging needs, and capture researcher-level outcomes.

- To help track the outputs of NAIRR-supported research, NAIRR management should provide standardized acknowledgement language that all NAIRR users are required to use when disseminating the results of their research through publications, websites, etc.

Performance/Success Indicators

	Resources	Innovation	Diversity	Capacity	Ethics
Inputs: Resource Investments	**Computational** **Data** **Testbed** **Training**	Number/volume/cost of resources provided	Accessibility of resources provided	Scalability of resources provided	Number of resources provided to promote responsible AI approaches
		Quality of resources provided			
		Diversity of resources provided			
Activities: Usage of Resources	**Computational**	Number of successful jobs		Average percentage of resource utilization	
		Average task turnaround			
	Data	Number of datasets accessed and reused			
		Number of new datasets contributed and subsequently reused			
	Testbed	Usage statistics for various testbeds			
	Training	Usage statistics for the various training courses and tools		User satisfaction ratings after completing training offerings	Percentage of users accessing training related to responsible AI research
				Percentage of users accessing training resources	
	Across All Resources	Domains/focus areas of research supported through the NAIRR	Number and diversity (incl. demographics, organizational affiliations, and geography) of individual and organizational users	Usage rates and character of resources accessed associated with different user types	
		Number and diversity of resource-provider partners		Educational level of users of the NAIRR	
		Usage rate of partner-provided resources		Number of proposals received, acceptance rate, and time to response	
		Primary disciplines/field of users of the NAIRR			

Performance/Success Indicators

	Resources	Innovation	Diversity	Capacity	Ethics
Outputs: Progress toward Goals	**Computational** **Data** **Testbed** **Training**	Number of research publications citing use of the NAIRR	Demographics of NAIRR users as compared to the demographics of the nation	Longer-term career trajectories of student users of the NAIRR	Impact metrics of research using NAIRR resources to advance approaches to responsible and trustworthy AI
		Number of papers citing use of the NAIRR selected for conferences		Use of the NAIRR resources in support of educational offerings and courses/tutorials	Adoption levels of responsible and trustworthy AI processes and tools developed through the NAIRR
		Impact metrics of publications from NAIRR-supported research		Number of students trained on NAIRR resources	Number of research publications related to ethical/responsible AI citing use of the NAIRR
		Number of patents developed by NAIRR users			
		Number (and download/use/adoption/citation rates) of new datasets, models, methods, and tools developed using the NAIRR			
		Capital raised/revenue of startups supported through the NAIRR			
Impact: Strategic Effects	**Computational** **Data** **Testbed** **Training**	Number of AI research papers published by U.S.-based researchers	Overall demographics of "AI R&D community" over time	Number of students graduating with AI degrees	
		Shifts in H-index of U.S. AI publications		Size of the national AI workforce	
		Improvement in efficiency and compute requirements of algorithms			

Interim Report: Outline & Process

LYNNE PARKER, DIRECTOR, NATIONAL AI INITIATIVE OFFICE, WHITE HOUSE OFFICE OF SCIENCE AND TECHNOLOGY POLICY

Proposed Outline

1. Introduction
   A. Unlocking the Potential of AI
   B. Charge to the NAIRR Task Force
   C. TF Approach to developing a NAIRR Roadmap
   D. A Vision for the NAIRR
2. Goals and Metrics for Establishment and Sustainment of a NAIRR
3. NAIRR Ownership, Administration, Governance and Oversight
4. NAIRR Resource Elements and Capabilities
   A. Data Resources
   B. Compute Resources
   C. Testbeds
   D. Educational Tools and Services
   E. User Interface
5. System Security and User Access Controls
   A. Security Requirements
   B. A Framework for Access Control
6. Safeguarding Privacy, Civil Rights, and Civil Liberties
7. Funding and Sustainment
   A. Establishment
   B. Agency Roles and Responsibilities
   C. Sustainment
   D. Implementation Milestones
8. Next Steps for the Task Force

Appendix A. Task Force Membership
Appendix B. Definitions
Appendix C. Briefers to the Task Force
Appendix D. Experts Engaged by NAIRR TF Working Groups
Appendix E. Public Input Provided in Response to Federal Request for Information
References
Abbreviations

Proposed Section Structure

1. Context on the importance of a specific resource element or system attribute for meeting the vision of the NAIRR
2. NAIRR TF/working group approach to researching needs and options for the resource element/attribute (as appropriate/relevant)
3. Key evidence in support of findings
4. Key findings about needs and options for the resource element or system attribute, with supporting evidence (explicitly enumerated, potentially interspersed throughout part 3)
5. Recommendations (explicitly enumerated) for design of the resource element or system attribute to meet identified needs, along with explanation for each recommendation

Timeline

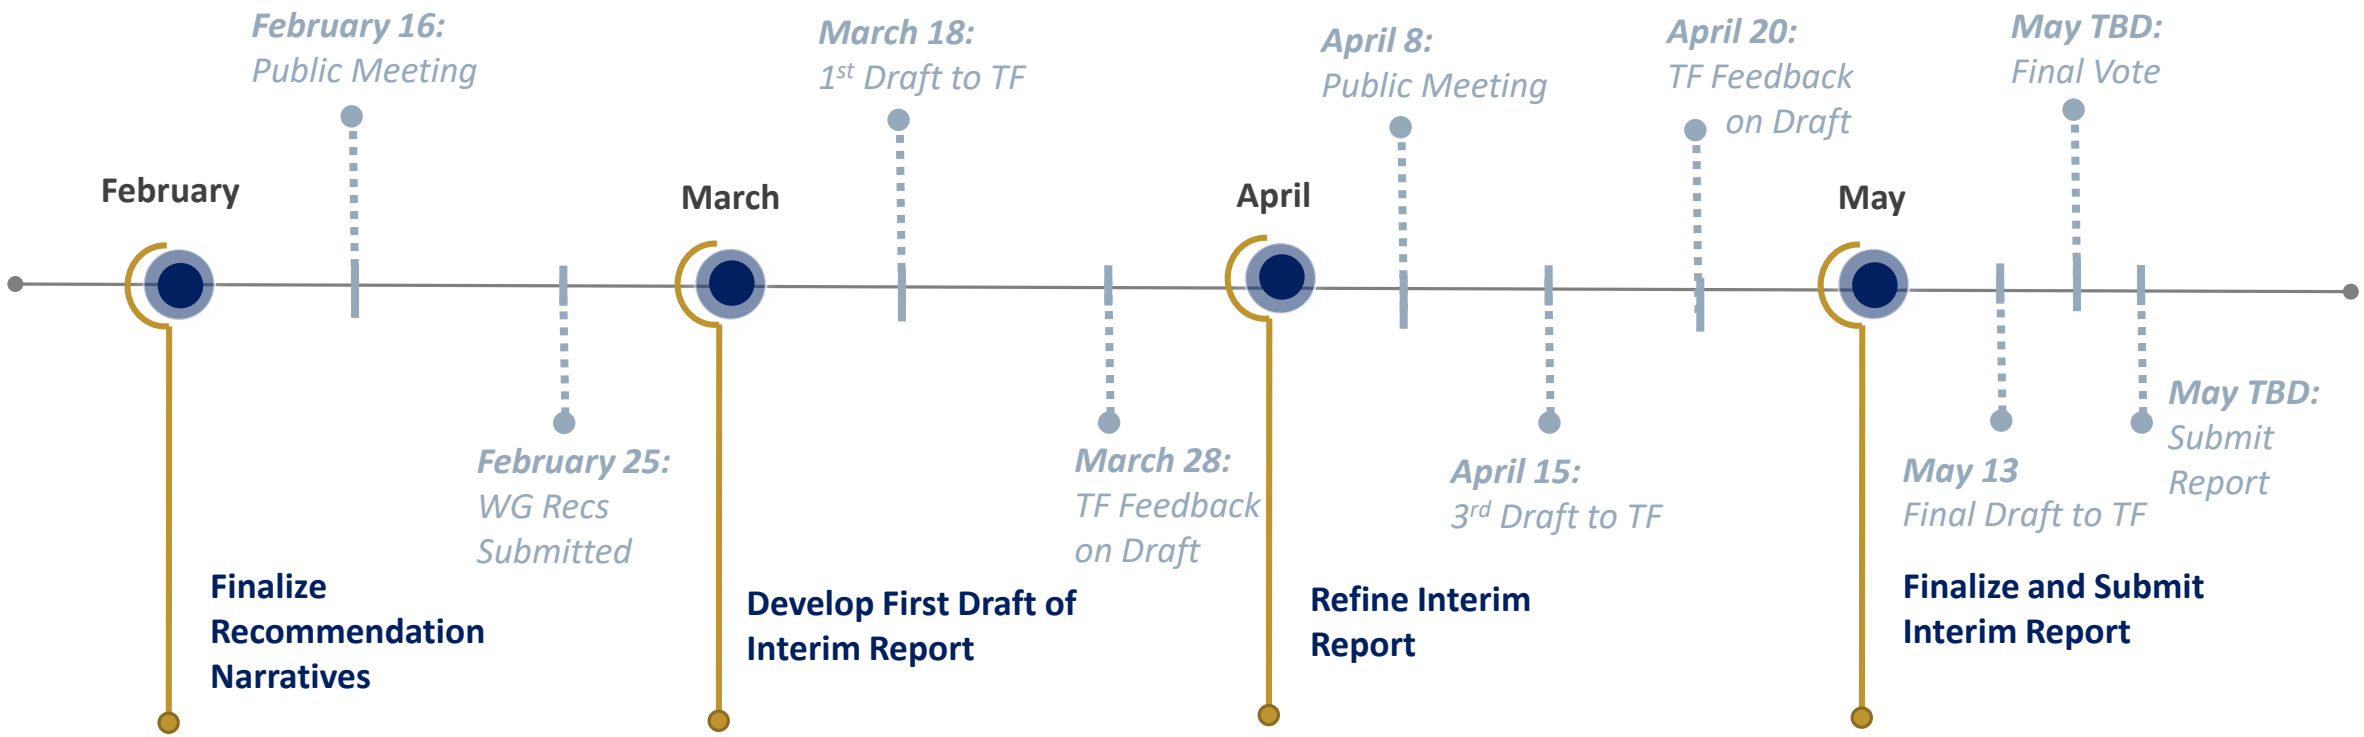
February
March
April
May
February 16:
Public Meeting
March 18:
1st Draft to TF
April 8:
Public Meeting
April 20:
TF Feedback on Draft
May TBD:
Final Vote
February 25:
WG Recs Submitted
March 28:
TF Feedback on Draft
April 15:
3rd Draft to TF
May 13
Final Draft to TF
May TBD:
Submit Report
Finalize Recommendation Narratives
Develop First Draft of Interim Report
Refine Interim Report
Finalize and Submit Interim Report