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

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

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

Attachment A: Aerospace Response and input to RFI questions

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

   A. **Goals for establishment and sustainment of a National Artificial Intelligence Research Resource and metrics for success**; The goals and metrics for this endeavor should state measurable, to the fullest extent possible, strategic end states that provide clarity and substance to outcomes. Characteristics of successful metrics are described below:

   - Metrics are key to
     - *Evaluating decision points*
     - *Making a defensible argument*
   - There are two different families of metrics
     - *Model Metrics*
       - Evaluating a machine model is a key aspect of determining whether it is useful in answering key business and operational questions. In this regard, this regard there is a fundamental tradeoff in performance between:
         - *Evaluation—how accurate a model is making inference (based on training and validation data)*
         - *Generalizability—how generalizable or scalable a model is (based on hold out data or ground truth)*
         - *Benchmarking—how well a model compares to other machine learning models*
       - Evaluation is intimately tied to the type of machine learning problem you are trying to solve
     - *Enterprise Metrics*
       - In order to properly track usage, derive effective insight on product usage, and deliver enterprise value, it is imperative to develop appropriate business-level metrics
       - There are two flavors of enterprise-level metrics:
         - *User Metrics: track service usage and provisioning behaviors*

- *Enterprise Metrics: track adoption across the national enterprise*

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 Research Resource; and

ii. A governance structure for the Research Resource, including oversight and decision-making authorities.

This plan should state clear roles and missions of Departments and Agencies including appropriations, authorities, and responsibilities up front to ensure maximum efficiency.

- *Emphasis of organization should focus on fostering innovation through sponsorship of cross-cutting capabilities*
  - Through coordination of funding
  - Making available large repositories for model training
  - Encourage sharing of information, data, code, best practices through online digital collaboration platforms, conferences, and representation on standards bodies
- *Aperture opens to include AI/ML within the Federal enterprise, open source, commercial, and academia*
- *Imperative to partner across USG and look for on-ramps*

C. A model for governance and oversight to establish strategic direction, make programmatic decisions, and manage the allocation of resources.

Governance and Oversight structure and process should be specified in writing to establish strategic direction, make programmatic decisions, and manage the allocation of resources. Conceptual models and processes for the National Artificial Intelligence Research Resource (NAIRR) are depicted in the two figures below.

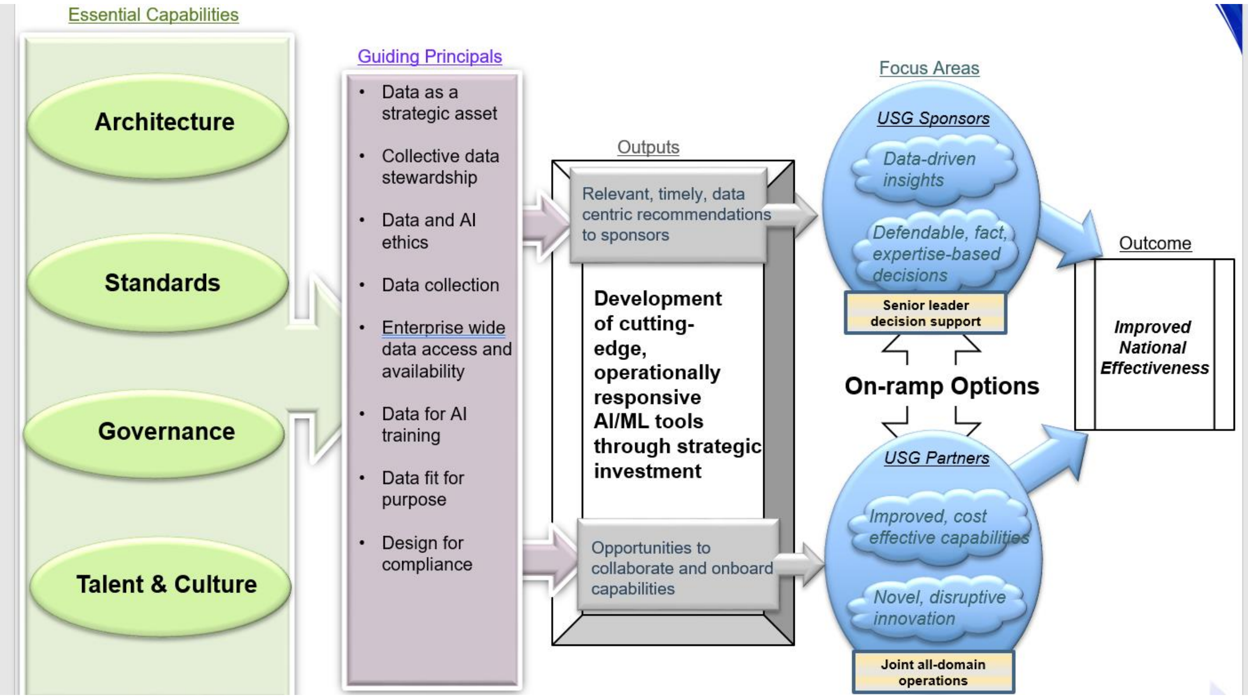
Essential Capabilities
Architecture
Standards
Governance
Talent & Culture
Guiding Principals
Data as a strategic asset
Collective data stewardship
Data and AI ethics
Data collection
Enterprise wide data access and availability
Data for AI training
Data fit for purpose
Design for compliance
Outputs
Relevant, timely, data centric recommendations to sponsors
Development of cutting-edge, operationally responsive AI/ML tools through strategic investment
Opportunities to collaborate and onboard capabilities
Focus Areas
USG Sponsors
Data-driven insights
Defendable, fact, expertise-based decisions
Senior leader decision support
On-ramp Options
USG Partners
Improved, cost effective capabilities
Novel, disruptive innovation
Joint all-domain operations
Outcome
Improved National Effectiveness

Figure 1: Holistic flow of inputs and decisions

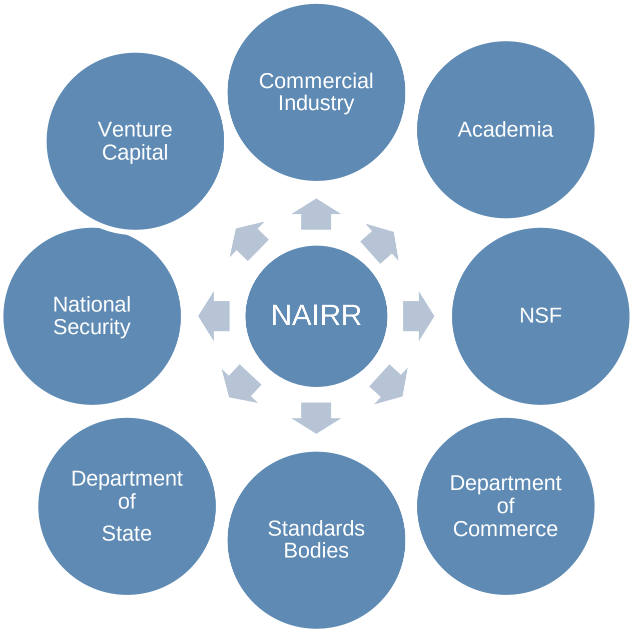

Figure 2: Partnership Ecosystem

NAIRR will:

- Support Academia in development of long-term AI/ML capabilities as well as scan the horizon for potential technology disruption
- Partner with NSF to identify, evaluate, and support funding opportunities for AI/ML research in the national interest
- Work with Department of Commerce to develop incentivization schema to foster US AI/ML business growth
- Collaborate with standards bodies to create systematic and repeatable metrics to evaluate AI/ML performance
- Advise and work with department of state to "hit the sweet spot" between specifying ITAR restrictions necessary to safeguard technologies critical to the national interest and cross-national collaboration
- Ensure that national security requirements are met through investment and premonition of AI/ML investments and technology across federal, commercial, and academic realms
- Partner with Venture capital to identify and share investment options in develop of organic AI/ML startups

D. **Capabilities required to create and maintain a shared computing infrastructure to facilitate access to advance computing resources for researchers across the country, including provision of curated data sets, compute resources, educational tools and services, a user-interface portal, secure access control, resident expertise, and scalability of such infrastructure.** Considerations of this characteristic would consist of, but not be limited to:

- Different combinations of data produce different analytics products
- The second most labor-intensive aspect of AI/ML development is curating data
- There are several nuances involved in data cleaning, sanitation, formatting etc.
- Having pre curated data is a boon for developers
- Linking these datasets (curated) to specific AI/ML problem areas and use cases would greatly facilitate knowledge discovery and make it efficient for developers to find what they are looking for
- Few, if any, analytics vendors keep their own infrastructure
    - *Cost is too high (especially for start-ups trying to hit VC backed revenue goals)*
    - *Not scalable to the really hard problems*
- Having a Graphics Processing Unit (GPU) enabled cloud infrastructure with persistence storage provided by NAIRR would be extremely helpful to developers as it would:
    - Give the ability to train models of appreciable complexity
    - Allow developers to devote funds to other activities (such as hiring researchers or funding other projects)
- In terms of infrastructure, providing support for containerization and orchestration is a key enabler of AI/ML DevOps

- It also ensures stability and consistency of deployment
- Providing a platform to share code, allow comments and exchange between developers, and promotion of a leaderboard across the variety of different AI/ML problems would greatly facilitate collaboration and also provide a resource from which developers could learn and apply to new problems
- Proper security access controls ensure that the NAIRR resource cannot be manipulated or otherwise repurposed to support the aims of malevolent actors

E: An assessment of, and recommended solutions to, barriers to the dissemination and use of high-quality government data sets as part of the NAIRR.

- There are numerous case studies/publications citing why AI products fail
- 85% fail due a handful of reasons (as shown below):

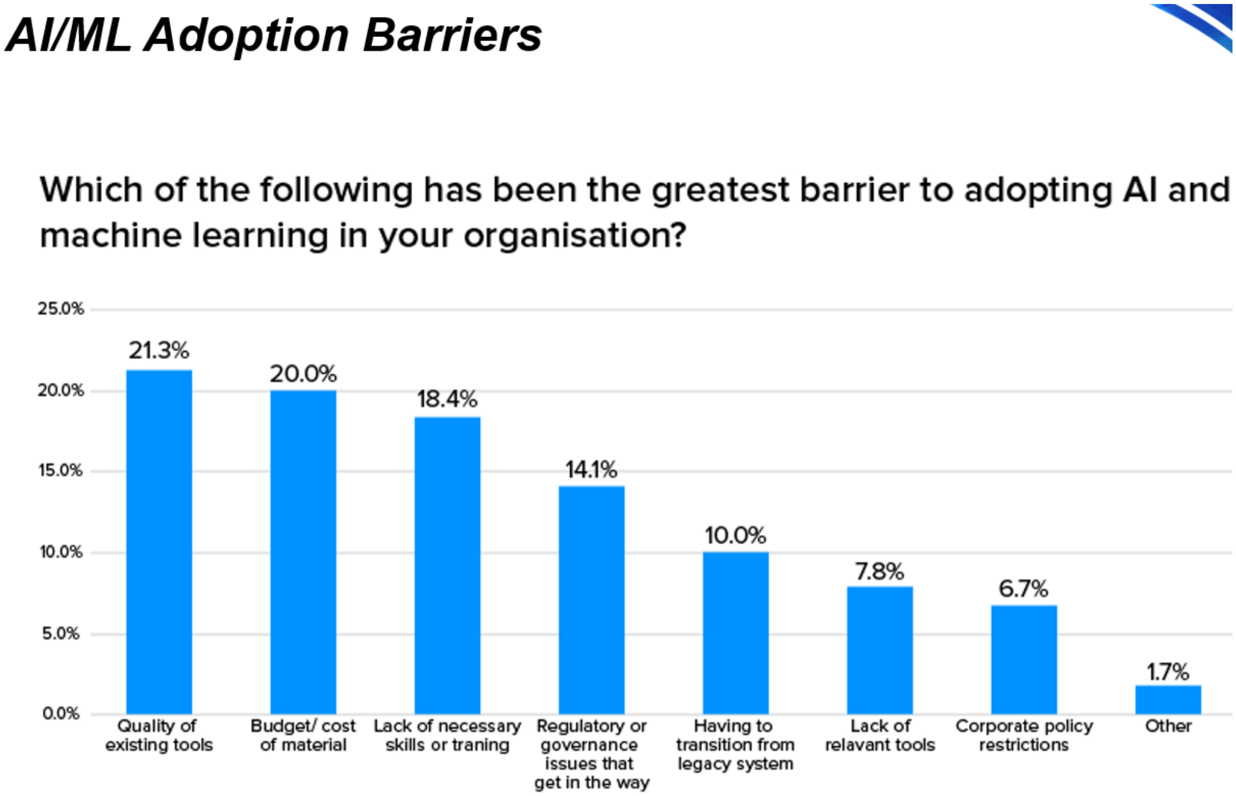

Figure 3: Barriers to AI Adoption

AI/ML Challenges

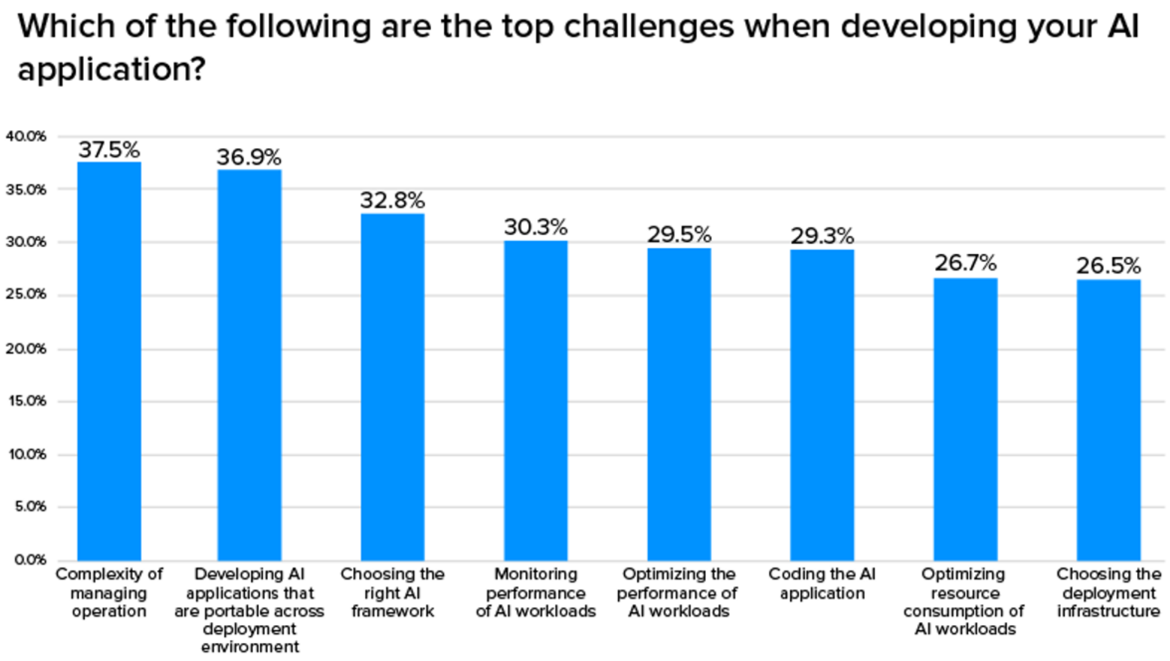

Which of the following are the top challenges when developing your AI application?

Source: Gartner 2020 AI Review

Figure 4: Top challenges to developing AI applications

Other Risks/Challenges:

- Data Science initial models don’t scale or are too experimental to be used by internal or external customers.
- Data science/ML models while brilliant and innovative don’t meet the business requirements or are too fragile to respond to change in the supporting data.
- AI initiatives are driven by the company’s internal IT organizations and inherit “waterfall” challenges.
- Companies don’t have the patience for the time it takes to deliver on AI/ML projects.
- A lack of a “Product” approach to AI/ML projects is core to project failure and increased risk.

Mitigation to these risks:

- Anticipating and planning against such risks only serve to strengthen the national interest

- It also helps provide insights as to when it's appropriate to either onboard or offboard AI/ML investments or otherwise course correct policy
- Once risks are identified, the must be addressed, though, some level of risk is necessary to innovate
- The question becomes how to balance risk versus reward
- Such guidance is key to building a roadmap

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

Cybersecurity, including access controls are an imperative and should be addressed up front. Based on extensive DoD Cybersecurity score carding, a significant percentage of incidents are due to Two Factor Authentication, Phishing, and Insider-enabled events.

Basic Cyber Hygiene to address these attack vectors is vital to maintaining the security and resilience of this initiative, which, if there is an attack, will receive significant attention. National resources in the Intelligence Community, the Department of Defense, and Federal Law Enforcement should be brought to bear for overall strategic overwatch and defense against the small, but significant portion of threat events enabled by foreign nation state capabilities.

G: An assessment of privacy and civil rights and civil liberties requirements associated with the National Artificial Intelligence Research Resource and its research.

Privacy and civil rights and civil liberties are foundational to this effort and should receive primary leadership and staff attention to ensure they are holistically included in the entire NAIRR effort.

H: A plan for sustaining the National Artificial Intelligence Research Resource, including through Federal funding and partnerships with the private sector; A clear, multi-year research and development investment profile, itemized by major program elements and further apportioned by Department and Agency is cross cutting and essential for effective governance.

- Identification and pursuit of funding opportunities and partnerships is critical to the NAIRR sustainment
- Because of funding limitations, the sheer complexity of AI/ML challenges, and the pace in which the landscape is changing no organization can provide a national AI/ML advantage if acting in isolation
- A partnership is an opportunity for NAIRR to align with another USG agency, company, or investment firm to support the development or deployment (or sometimes both) of a critical AI/ML capability

- Partnering improves both potential technology reach and joint operational effectiveness, but it also introduces a few complexities for development and deployment
- There are many different types of funding & partnering options, each with associated tradeoffs
- This is not a one size fits all.
- Different options must be exercised based on circumstance
- The key is conducting the requisite due diligence to find which option is appropriate (to mitigate risk and ensure an optimal outcome) and then structure a portfolio of these options in order to promote national AI/ML interests

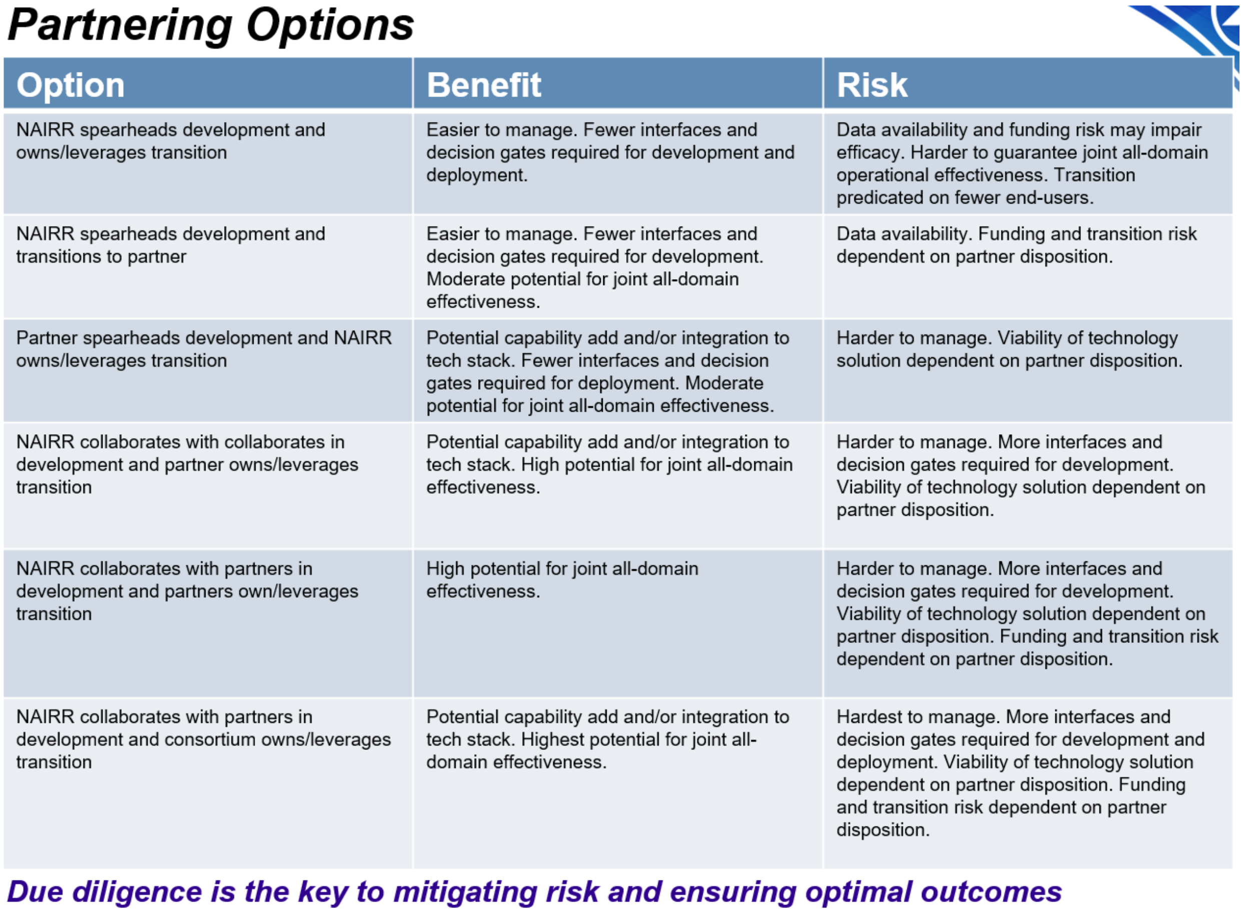

Partnering Options

Option	Benefit	Risk
NAIRR spearheads development and owns/leverages transition	Easier to manage. Fewer interfaces and decision gates required for development and deployment.	Data availability and funding risk may impair efficacy. Harder to guarantee joint all-domain operational effectiveness. Transition predicated on fewer end-users.
NAIRR spearheads development and transitions to partner	Easier to manage. Fewer interfaces and decision gates required for development. Moderate potential for joint all-domain effectiveness.	Data availability. Funding and transition risk dependent on partner disposition.
Partner spearheads development and NAIRR owns/leverages transition	Potential capability add and/or integration to tech stack. Fewer interfaces and decision gates required for deployment. Moderate potential for joint all-domain effectiveness.	Harder to manage. Viability of technology solution dependent on partner disposition.
NAIRR collaborates with collaborates in development and partner owns/leverages transition	Potential capability add and/or integration to tech stack. High potential for joint all-domain effectiveness.	Harder to manage. More interfaces and decision gates required for development. Viability of technology solution dependent on partner disposition.
NAIRR collaborates with partners in development and partners own/leverages transition	High potential for joint all-domain effectiveness.	Harder to manage. More interfaces and decision gates required for development. Viability of technology solution dependent on partner disposition. Funding and transition risk dependent on partner disposition.
NAIRR collaborates with partners in development and consortium owns/leverages transition	Potential capability add and/or integration to tech stack. Highest potential for joint all-domain effectiveness.	Hardest to manage. More interfaces and decision gates required for development and deployment. Viability of technology solution dependent on partner disposition. Funding and transition risk dependent on partner disposition.

Due diligence is the key to mitigating risk and ensuring optimal outcomes

Figure 5: NAIRR Partnering Opportunities

I: Parameters for the establishment and sustainment of the National Artificial Intelligence Research Resource, including agency roles and responsibilities.

- Agency Roles and Responsibilities addressed in (A).
- What's important for (I) is that NAIRR must "own" the due diligence process, determine which partnering/investment options are appropriate, and then structure a portfolio of options to serve the national AI/ML interest.

- Aerospace suggests the following due diligence process.

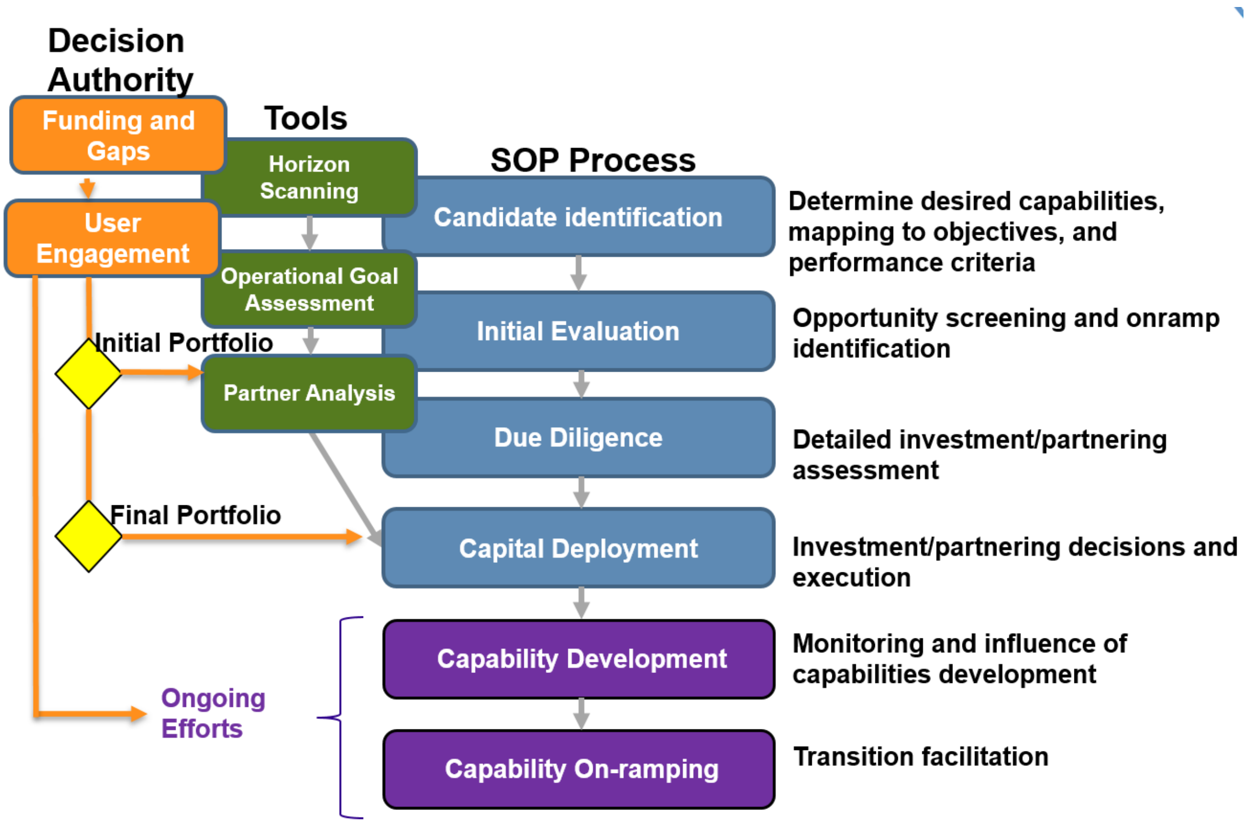
Decision Authority
Funding and Gaps
User Engagement
Tools
Horizon Scanning
Operational Goal Assessment
Partner Analysis
Initial Portfolio
Final Portfolio
Ongoing Efforts
SOP Process
Candidate identification
Initial Evaluation
Due Diligence
Capital Deployment
Capability Development
Capability On-ramping
Determine desired capabilities, mapping to objectives, and performance criteria
Opportunity screening and onramp identification
Detailed investment/partnering assessment
Investment/partnering decisions and execution
Monitoring and influence of capabilities development
Transition facilitation

Figure 6: NAIRR Due Diligence Process

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

Aerospace recommends prioritizing in this order: B, A, C, D, with the other items being specific annexes and sub-initiatives to C.

Furthermore, security access controls are considered a prerequisite for any front facing platform deployment. Following this, in order of precedence:

- Data Repository
- Data Curation
- Infrastructure
- Code repository
- Developer Engagement Platform
- Leaderboard
- Onboarding of additional learning resources

Critical to this is the development and integration of enterprise metrics in order to understand what elements of the NAIRR platform are most used, under-used, etc. A/B testing and user outreach should be integrated into DevOPs so that features most useful to end-users are offered.

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

   Ethical implications of AI include liability and law and are necessary conditions to integrate moral, societal, and legal values with technological developments in AI. Responsible AI is a priority and this theme and responsibility should be a foundational part of Item C.

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

   Broad existing initiatives exist in the following areas that should be leveraged to the maximum extent possible:

   - Quantum and High-Speed Computing
   - Cybersecurity – Specifically the Comprehensive National Cybersecurity Initiative (CNCI) from 2007 – 2014. 12 Initiatives plus Enablers – strong model for this AI initiative
   - All recent NIST Special Publications on AI, SP 1270 Identifying and Managing Bias in Artificial Intelligence
   - Advanced Autonomy
   - Private Sector Big Data Analytics
   - Advanced Manufacturing
   - Trusted Micro Electronics
   - Innovation (for example, the Department of Defense's Developmental Innovation Unit (DIU) in Mountain View, CA)

   Several building blocks exist that can be leveraged. In terms of company selection for partnering, both Crunchbase and Pitchbook maintain voluminous stores of information that can be access either online or via API call. The Microsoft Academic Graph has data on over 500 million academic publications that can be mined for the purposed of technology identification, horizon scanning, or research collaboration.

   IEEE and ACM have active AI/ML panels that could be tapped into to provide situational awareness. NSF could also be tapped into as well. Many academic AI/ML institutions almost have incubators/angel investment arms that could help provide targeted engagement for investment sharing. In-Q-Tel and DIUX could likewise provide investment intelligence from the standpoint of alignment with national security AI/ML requirements.