PUBLIC SUBMISSION

Received: May 14, 2025
Tracking No. mao-d6x2-1rkr
Comments Due: May 28, 2025 **Submission Type:** API

Docket: NSF-2025-OGC-0001
NITRD_FRDOC_0001

Comment On: NSF-2025-OGC-0001-0001
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

Document: NSF-2025-OGC-0001-DRAFT-0076
Comment on FR Doc # 2025-07332

Submitter Information

Organization: Kaiser Research LLC

General Comment

Kaiser Research respectfully submits our response via the attached PDF.

Attachments

KZR-NSF-2025-OGC-0001

Kaiser Research

Response to the 2025 National Artificial Intelligence Research and Development Strategic Plan RFI

Submitted by Bill Patrowicz, Principal, Kaiser Research LLC

Docket ID No. NSF-2025-OGC-0001

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

Overview

Kaiser Research is a strategic, non-technical partner focused on **strategic computing** — the intersection of high-performance computing (HPC), artificial intelligence (AI), and quantum technologies with real-world, mission-driven challenges. We work with national laboratories, FFRDCs, mission-aligned industry, and public-private ventures to turn advanced computing capacity into durable national capability.

We write from a practical vantage point: close to the compute, embedded in institutions, and engaged in efforts that link infrastructure, innovation, and national interest. This submission focuses on a persistent gap in AI strategy: the need to build not just capabilities, platforms, and infrastructure — but **ventures and institutions** — capable of delivering long-horizon, public-purpose outcomes through computing.

Reframing The Challenge

The 2023 National AI R&D Strategic Plan rightly emphasizes long-term investment, safe and trustworthy systems, and public-private collaboration. But between these aims and durable outcomes lies a quieter, structural question: **how do we organize ourselves to deliver at national scale?**

Today's serious conversations about AI infrastructure cluster around five poles:

- **Hyperscalers** (AWS, Azure, GCP), whose domination of AI workloads raises questions of sovereignty, lock-in, and public utility;
- **Public initiatives** like the EuroHPC JU in Europe and NSCI, focused on equitable access and national capacity;
- **Commercial HPC-as-a-Service firms** (Rescale, CoreWeave, Lambda Labs), who monetize access without creating enduring infrastructure;
- **Discovery bottlenecks** — in fusion, climate, and biomedicine — where compute constraints stall scientific progress;
- And **venture-backed "AI foundries"**, where compute, capital, IP, and talent are bundled for rapid product acceleration.

What's missing is an operational model that blends **governance, translational depth, and institutional trust** — one that isn't just scalable, but mission-aligned. That's where U.S. leadership is still possible.

A Model Worth Building Toward

We often describe the needed structure as a **Mayo Clinic for computing** — an institution defined not by capacity, but **by a network platform for comprehensive, integrated, mission-oriented compute** that is sought after to address the world's most complicated, interdisciplinary challenges.

Mayo is a durable, trusted vessel for complexity. Patients enter not for individual services, but for coordinated, long-horizon care. It's built for hard problems.

Translate that to computing:

- Where compute scientists, ML engineers, and domain experts operate as **a mission-driven team**;
- Where **infrastructure is embedded in the institution**, not rented;
- Where value comes from solving **high-stakes scientific and industrial problems**, not monetizing GPU hours;
- And where long-term initiatives — from space to energy — are not special projects, but the organizational core.

This model diverges from both hyperscalers and commodity HPC. It prioritizes **execution over access, continuity over speed, and outcome over infrastructure metrics**. In the use-cases alluded, it will also deliver less cost and more value.

We call this model **The Mission Venture**. Others may be inclined to call it a public-private venture (PPV), the successor to the PPP, designed for multi-lateral co-investment, shared governance, shared assets, and strategic continuity.

Strategic Alignment

This vision aligns directly with the 2023 Strategic Plan priorities:

- **Priority 1:** Long-term investments require institutional homes — not pilot programs;
- **Priority 2:** Human-AI collaboration thrives when interdisciplinary teams are the rule, not the workaround;
- **Priority 3:** Trustworthy AI is more credible when it's governed inside **mission-aligned ventures**, not abstract markets;

- **Priority 5:** Shared compute becomes effective when embedded in institutions that deliver outcomes, not just access;
- **Priority 8:** Public-private partnerships evolve into **public-private ventures** — designed not to share cost, but to share purpose, infrastructure, and returns.

Just as the Department of Energy once built institutions around nuclear capability, this era demands new institutional forms around compute.

Closing Thought

AI does not operate in the abstract. It runs on silicon, in buildings, with people, governed by institutions. If we want lasting capability, we must build **ventures that can hold it** — not just more flexible APIs and technical know-how, but places where long-horizon computing delivers long-term public outcomes and can provide scalable supply to non-linear demand.

Mission Ventures and PPVs are one such model. Durable, co-invested, and mission-driven, they allow the U.S. to move beyond grants and service contracts — toward institutional structures that can execute at scale.

The challenge ahead is not just technical. It's institutional. And the opportunity isn't just to build systems — it's to build **the kinds of places where those systems can matter.**

