NSF Tech Labs RFI response:
The case for the Quantum Biology Institute

Core topics of our response. Novel, high-quality empirical data generation; cutting-edge instrumentation-driven platforms; release of intermediate, broadly reusable research outputs.

III. If you have feedback, questions, areas of concern, suggestions on aspects of the NSF Tech Labs, and/or other information you want to have considered as a part of this RFI, this is an opportunity to provide focused commentary.

It appears that the NSF Tech Labs program borrows key elements of the ARPA model to direct public funding toward high-risk/high-reward science, rather than ARPA's focus on technology.

technology: the use of knowledge to extend the powers of humankind

funding	low-risk/low-reward	high-risk/high-reward
public	BARDA, FNIH, etc.	ARPAs
private	'Big Pharma', etc.	startups

science: the use of empirical observation to extend the knowledge of humankind

funding	low-risk/low-reward	high-risk/high-reward
public	NSF, NIH, etc.	**NSF Tech Labs**
private	March of Dimes, etc.	HNWIs, 'DeSci'/crypto

Altering the ARPA model to become a model for breakthrough science: three science enterprise models. Recast the science challenge as a technology challenge by funding new ways to generate empirical data. We see three complementary approaches:

- **data-intensive platforms** (e.g., Human Genome Project; BRAIN Initiative; AI agents for science);
- **dissemination of low-tech, scalable platforms** (e.g., frugal science; citizen-science efforts);
- **cutting-edge instrumentation-driven platforms**, developed first in specialized laboratories, and requiring multi-year engineering (e.g., PCR; lattice light-sheet microscopy; efforts at the Quantum Biology Institute).

Metrics will be the crucial design choice for the Tech Labs program. First, metrics for science should differ from those for technology: technology can often be judged by whether it works, whereas science should be judged by whether it yields rigorous, informative results — whether confirming or refuting a hypothesis — that then necessarily generate new science questions: the experimental path to technology resolves scientific roadblocks along the way. Second, metrics work for iterative research because one already knows what good, anticipated data look like. In contrast, metrics do not work well with breakthrough research because one does not know what successful results will look like, as they are likely unprecedented. Yet metrics play an important role in ARPA-style programs: they provide accountability, enable course-correction, and help distinguish progress from busywork.

Metrics should focus on novel, high-quality empirical data generation to achieve targeted outcomes. For each of the three science enterprise models outlined above, work will necessarily begin at low TRLs; accordingly, for low TRLs, overly precise metrics are unrealistic. Instead, the key Tech Labs metric should be the extent to which novel, high-quality empirical data are being generated to directly support implementation. Such data should be reliable irrespective of whether they support or refute a hypothesis, incorporating negative results to invalidate claims, and enabling unbiased exploration of the search space. As we argue repeatedly below, requiring the release of intermediate, broadly reusable research outputs can provide a clear view into whether the above-defined metric is being met.

2. What, if any, substantive comparative advantage (as compared to standard grants and other existing NSF programs) could the NSF Tech Labs program model provide in efforts to accelerate and advance U.S. competitiveness — either across various key technologies or within a specific technology focus area?

#1. Filling the existing career-path gap for execution-focused professional scientists. More research scientists are trained at legacy research institutions than there are stable research positions within them. The NSF Tech Labs program can provide an alternative, stable career track for execution-focused professional scientists to pursue foundational R&D outside the university PI/trainee model. By funding teams with minimal teaching and training obligations, this program can: emphasize sustained technical execution over academic bureaucracy; retain expertise through long-term, well-compensated roles (with appropriate at-will flexibility); and preserve institutional knowledge that is otherwise lost in year-by-year temporary academic appointments.

#2. Enabling nimble, flexible teams. Research labor at universities is organized around training and tenure, which dilutes execution, slows staffing changes, and makes it hard to align personnel with fast-changing technical needs. By funding teams of professional scientists in at-will roles (rather than tenured PIs and fast-rotating trainees), Tech Labs can hire for specific expertise, adjust headcount quickly, and maintain performance accountability as projects evolve — while sustaining core capabilities and institutional knowledge over the grant's long time horizon.

#3. Fostering true interdisciplinarity. Standard grants can support interdisciplinary work, but PI incentives, departmental structures and silos, ever-changing and competing interests for tenure lines, and trainee-based labor at universities often reduce such work to 'collaboration on paper.' By funding integrated teams of professional scientists, Tech Labs can hire and retain expertise across fields, change team composition as technical needs evolve, and sustain de facto, deep cross-disciplinary integration — exactly what frontier areas (e.g., quantum biology) require to produce breakthrough technologies and strengthen U.S. competitiveness.

#4. Supporting mission-oriented teams with clear deliverables and technology outputs. Tech Labs can fund mission-oriented teams organized around defined objectives rather than individual PI agendas. This structure supports time-bounded efforts with concrete technology outputs, such as referenced implementation models, validated datasets, testbeds, contributions to standards, and robust open software — all of which are often poorly aligned with university incentive structures. Importantly, projects should retain a substantial basic-science component (since new science drives new technology) while maintaining a clear roadmap for translating that science into scalable, economically valuable technologies. Because Tech Labs careers need not be driven primarily by publications, the program can explicitly reward the dissemination of these outputs, accelerating translation and adoption.

4. What program design choices would contribute to the success of the Tech Labs mechanism?

Privileging interdisciplinary areas that integrate engineering and basic science above all else. Established industries tend to look like self-contained 'subject areas' precisely because they are already mature: their methods, problems, and talent pipelines have been consolidated into a single recognizable domain. By contrast, breakthroughs that create new industries usually emerge at the seams between fields and therefore appear, at the time of their inception, too interdisciplinary to fit neatly into any one bucket. Tech Labs topics should thus be treated as loose targets rather than strict categories that proposals must conform to. Indeed, a proposal that fits cleanly inside a single area, without substantial work outside it, should be viewed skeptically as potentially insufficiently novel to catalyze a new industry. In other words, topics should not silo truly interdisciplinary projects. Consider PCR: when Mullis introduced it in 1983, it drew on what would then have been seen as disparate areas: molecular biology, biochemistry, biophysics, cloning and genetic engineering, self-assembly...; yet, today these are comfortably grouped under the established umbrella of 'biotechnology.'

Daring to reduce the traditional NSF focus on training. It is tempting to equate 'new career paths' with expanded student support; however, Tech Labs should minimize teaching and student training so core teams can focus on mission-critical R&D and avoid spending on nonessential activities. While internships can be valuable, making student involvement a program requirement would tether Tech Labs to universities and risk capture by traditional academe, since students are inseparable from the institutions that enroll and credential them, especially at the graduate level. Instead, career-path goals should target trainees who are not full-time students via time-bounded roles (e.g., full-time internships and one-year fellowships) and via credentials such as certifications or professional doctorates rather than M.S. or Ph.D. degrees.

Daring to avoid prestige metrics. Because Tech Labs is designed to build execution-focused, mission-oriented teams, selection and evaluation should not hinge on academic prestige signals (institutional rankings, celebrity PIs, h-index, or journal venues). Those proxies reward past academic status and tend to concentrate resources in a small set of elite researchers, recreating the very PI-centric system Tech Labs is meant to complement. Instead, selection should emphasize execution — ideally already demonstrated outside the shelter of traditional academia: the ability to secure resources, assemble and lead interdisciplinary teams, deliver on ambitious technical milestones, and articulate a credible vision for sustained R&D and technology translation. This shift broadens the talent pool, enabling outstanding engineers and scientists from industry, national labs, startups, and less-heralded institutions to lead; in this way, Tech Labs will not simply be staffed by the same small cadre of senior academics, but by the teams most likely to execute and translate high-risk ideas into durable technological platforms.

4.a. Which types of teams and organizations should be considered eligible to apply for the NSF Tech Labs program? What restrictions on team eligibility should be in place to maximize speed and ensure novel impact?

Prioritize independent, new institutions and avoid recapture by legacy academia. Tech Labs should favor independent, newly formed startups and purpose-built research institutes, and discourage proposals anchored in legacy universities or large incumbent organizations that already have substantial NSF support and for which a Tech Labs award would be less transformative. The point is not to steer additional money to established structures, but to test a genuinely different model of innovation. Faculty and students need not be barred outright, but Tech Labs should be structured as a full-time, execution-driven enterprise — closer to a technology startup than a PI-led academic lab. Accordingly, the team lead (and a substantial fraction of the team, e.g., at least half) should be full-time professionals without primary obligations as tenured faculty, full-time teachers, or full-time students. Otherwise, Tech Labs risks becoming another funding stream for the same people and incentives that traditional NSF grants already support. Finally, independence from legacy academia must be real: 'talent mobility pathways' and 'infrastructure subcontracts' should not become backdoors for routing funds to universities. If teams can subcontract core work back to academic labs or maintain extensive academic appointments while executing the project, then Tech Labs is effectively funding university research with extra steps, less oversight, and a weaker alignment to its mission. Tech Labs should also be cautious about treating an academic leave of absence as sufficient commitment. Leading a Tech Lab is a multi-year, full-time undertaking that is difficult to reconcile with maintaining an active PI lab: research groups cannot simply be 'parked' for 5 to 7 years without substantial wind-down and reorganization, and a nominal leave may keep the team lead structurally and professionally tethered to the university. Moreover, Tech Labs work may not optimize for academic success metrics (e.g., frequent publications), creating misaligned incentives for leaders who expect to return to academia. In practice, a Tech Lab that succeeds will create larger operational and translational responsibilities — products, partnerships, scale-up, and follow-on funding — that make a return to academia unlikely. For these reasons, the program should prioritize leaders and teams able to commit fully from the outset, with governance and incentives aligned to Tech Labs outcomes rather than academic continuity.

4.b. Is the proposed timeline for Phase 0 (9 months), Phase 1 (24 months), and Phase 2 (24+ months) well-calibrated to support the program's strategic objective of achieving high impact, accelerated outcomes? If not, what adjustments should be made and why?

Initial fund disbursement must be timely. Teams may require rapid early funding for upfront investments, especially for enabling infrastructure and other critical capabilities needed to execute at scale. It is therefore crucial that teams be able to deploy \$10–50M per year by Phase 1.

Feedback on Phase 0. Phase 0 gives teams time to form and refine milestones before committing fully, lowering the barrier for non-traditional applicants without institutional backing, which we appreciate.

Feedback on Phase 1. Phase 1 milestones should explicitly permit learning-driven pivots: rigid deliverables set nine months earlier will penalize teams that uncover unexpected results, even though that is often where breakthroughs emerge. Longer Phase 1 time horizons (~36-48 months) would help to account for the uncertainty in dealing with scientific barriers towards advancing TRLs, thus increasing the likelihood that genuinely transformative technologies mature to the point of adoption; this is especially true for the scientific enterprise model that relies on cutting-edge instrumentation-driven platforms, as argued above.

Feedback on Phase 2 and beyond. NSF should define and pre-negotiate well in advance clear transition pathways for successful teams at the program's conclusion by engaging government, industry, philanthropic, and private-sector partners, so that high-performing teams can sustain momentum rather than disband when the award ends. Potential pathways include follow-on federal support, partnerships with FFRDCs and national laboratories, procurement routes for pilots/prototypes or deployment contracts, and private partnerships.

4.c. How should IP rights be structured to support maximum success and impact?

IP should be as unstructured as possible. Like successful startups, the most successful Tech Labs efforts will mature into new companies, products, and ultimately entire industries, possibly through partnerships or acquisition by larger industrial players. The program should therefore avoid imposing a one-size-fits-all licensing or IP regime that may seem sensible today but which will likely prove misaligned once the technical and market realities are known. Premature, rigid IP terms can depress valuation, deter follow-on investment, and lock teams into unnecessary constraints. Instead, Tech Labs should set only minimal, high-level guardrails and allow the emerging industries to develop appropriate IP models as they mature. Existing national-security and export-control frameworks already address the risk of adversarial capture; if those safeguards are insufficient, Tech Labs is not the right venue to redesign them.

Balance open infrastructure with protectable IP. Our team strongly supports open science while preserving the ability to protect and commercialize genuinely novel inventions. A useful precedent is NSF's own Prototype Open Knowledge Network (NSF 23-571), which explicitly frames open, shared data infrastructure as a public good. Tech Labs can adopt a similar posture by requiring the release of intermediate, broadly reusable research outputs (e.g., control or reference datasets, benchmarks and contributions to standards, documentation, foundational instrument designs, and robust open software), while allowing teams to retain patent rights and protect proprietary know-how needed to attract follow-on investment (e.g., IP-enabling molecules and datasets).

4.d. What degree of independence is optimal to ensure the flexibility, freedom, and speed required for the Tech Labs initiative?

The core grantmaking challenge is establishing trust between funder and scientist. Scientists may naturally argue for maximal autonomy, but we recognize the central tension in grantmaking: balancing operational independence with responsible stewardship of public funds. In frontier technical work, hands-on management is often counterproductive because only the technical team can reliably judge how to proceed as uncertainty unfolds. At the same time, technical excellence does not automatically translate into financial or operational discipline, especially when institutional constraints and incentives are misaligned. We argue below that trust should be earned through lean, high-signal reporting and transparency, and that, once earned, it should be rewarded with meaningful operational autonomy, with the expectation of ongoing dialogue.

Require lean, high-signal, transparent reporting. Tech Labs should replace frequent written reports with sparse, high-signal check-ins, preferably via in-person or video site visits, so teams spend more time on R&D and less time servicing bureaucracy. Formal milestones and funding checkpoints should be clear and agreed upon between the team and NSF at the outset of each Phase's launch; checkpoints should not exceed two or three per Phase. Continual interim reports, micro-benchmarks, and constant 'status' requests mainly reward teams that learn to game reporting rather than teams that make technical progress. Transparency can instead come from two lightweight mechanisms: **(i)** real-time financial visibility (e.g., monthly online budget/burn-rate reporting), and **(ii)** routine release, automated if possible, of intermediate, broadly reusable research outputs, as argued above. This gives NSF program managers ongoing, granular visibility into progress without imposing constant reporting burdens on scientists, while preserving a credible option to pivot or terminate funding when milestones are missed.

4.e. How should funding be allotted to each proposed Tech Labs? What factors — for example: team size, team expertise, infrastructure needs, growth trajectory — should NSF consider to determine appropriate funding amounts to support successful Tech Labs teams?

Budget for real infrastructure and reward demonstrated efficiency. Independent teams and new institutes often lack access to universities or legacy 'core facilities,' so credible plans (and budgets) for buying services or building in-house capacity should be treated as central to award sizing. Importantly, infrastructure need not imply enterprise pricing: efficient teams can deliver robust systems without large vendor contracts or overhead. In addition, salaries and benefits must be competitive to recruit and retain top talent.

Prioritize nimble, intentionally capped team sizes. In a R&D organization, there are natural tradeoffs between its size, the scale and complexity of its goals, and the agility of its operation. For this reason, Tech Labs awards should not be targeted at any particular absolute team size or budget range. Rather, they should target a favorable ratio of award size-to-outcome probabilities: smaller teams should not be penalized and larger teams should not be privileged. Still, one comparative advantage of Tech Labs is agility, so teams should resemble high-performing startups: small, tightly led, and able to move quickly. Teams should be kept deliberately compact (e.g., two highly engaged leaders with a core staff on the order of dozens, not hundreds); beyond that size, organizations typically require layers of middle management, which slow iteration and erode the very agility the model is meant to create. In addition, teams that physically co-locate to work together should be prioritized, especially for experimentally intensive projects; and pro forma, 'box-ticking' cross-sector collaborations should not be required.

5. What opportunities do you see for synergy with research and development efforts that are or could be funded by industry or philanthropic organizations? What partnership structure would allow Tech Labs to leverage federal and private support for maximum benefit?

Unlock underutilized industrial infrastructure. Industry partners could contribute in-kind access to specialized, intermittently used equipment and to otherwise idle compute capacity, lowering capital needs and accelerating execution for Tech Labs teams.

Partner with mission-aligned philanthropy to co-fund shared infrastructure. Philanthropic organizations interested in new models of science, especially those focused on open science and novel research infrastructure, could co-fund foundational platform work that benefits entire fields rather than the outputs of any single Tech Labs team.

6. What translational problems, challenges and/or bottlenecks could be addressed within 3-7 years with this program design?
3.i. What is your vision for substantially advancing the state of practice in your chosen topic of focus, and why is the Tech Labs funding mechanism a superior approach to pursuing that vision?

Quantum biology unlocks a new way to actuate on biology. If an unambiguous quantum-to-biology link can be established and then rationally leveraged through engineering, a wide variety of industries will be empowered: from theranostics, biomanufacturing, and space exploration via magnetosensing and actuation; to improved drug discovery and energy efficiency via quantum-enhanced models of molecular recognition. Moreover, harnessing how biology uses quantum phenomena in warm, noisy environments may yield design principles for more robust quantum technologies. Traditional funding models avoid paradigm-shifting fields such as this.

Tech Labs can build the quantum biology industry. A persistent barrier to unlocking the quantum biology industry is the absence of reliable, high-quality quantum measurements from biological samples that can serve as a 'codebook' for how to manipulate biology to achieve function. A dedicated Tech Lab such as the Quantum Biology Institute can close this gap within 3-7 years. The Institute, a FRO identified in Convergent Research's gap map, builds and deploys cutting-edge, instrumentation-driven platforms that generate bona fide quantum data from living systems. It also provides a dedicated home for quantum biology by selecting and executing a portfolio of the most promising research directions for scientific and technological impact. Headquartered in Los Angeles and organized as a 501(c)(3) non-profit, it is already pioneering the real-time release of intermediate research products (e.g., electronic lab notebooks and micropublications) under an Apache 2.0 License. IP-sensitive enabling molecules and datasets can be kept non-public at the team's discretion when controlled access is necessary to preserve downstream commercialization options.