2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

BREAKOUT SUMMARY · INCREASING VISIBILITY OF AWARDS & OUTPUTS

Visibility = findability, not publicity

NSF CI + WORKFORCE
CSSI · CyberTraining · SCIPE

Facilitator: Andreas Goetz · Reported by: kc · PI meeting for CSSI, CyberTraining, and SCIPE, 15 Sep 26

DIAGNOSIS

Producers cannot reach users who need their work; users cannot tell what exists. **Solve discovery and visibility follows.**

PROGRESS METRIC

% of awards with at least one resolvable, non-publication artifact on award record.

TOP RECOMMENDATIONS (more in summary https://go.osu.edu/nsf26)

1. **Extend NSF PAR to software, data, and curricula**
   Publications already on award pages — widening scope is highest-leverage change.
2. **Require a persistent identifier for award products**
   Accept any DOI-issuing repository; display compliance publicly.
3. **Fund domain-scale portals, not an uber-portal**
   People follow people they trust. Fund stewardship (e.g., stipends), not just content.
4. **Drive onboarding friction toward zero**
   Credential-free or pre-provisioned first runs; fund outreach.

ALREADY WORKING, AT SCALE:

~200
poster DOIs mandated for this meeting: one-click form + public MISSING flag

30+
universities feed ACT-CMS curriculum portal, sustained by fellowships

~21k
training hours on FABRIC's Artifact Manager — 1-click deposit

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

Sustainability of research software and CI

CSSI/SCIPE/CyberTraining PI meeting 2026
Presenter: Melissa Kline Struhl

The Problem

Software written by PhD students, who leave, or faculty members, with myriad other responsibilities

Academic incentives reward papers, not maintenance. When initial NSF funding ends, projects face a cliff. Institutional knowledge walks out the door. Successful tools fail. Researchers reinvent the same solutions repeatedly.

This is a waste.

The good news: Thanks to programs like CSSI, PESOSE, FAIROS, and more, **building and establishing** infrastructure has a much better chance!

The Cliff

CSSI funding is designed to end - Transition to Sustainability is designed to help projects identify and transition to “other avenues of support.”

Now what?

Everyone knows a telescope facility stops working if you don’t maintain it.

Everyone also knows that some (many?) telescopes can’t be funded as for-profit ventures - commercialization pathways don’t fit all projects.

Successful Examples - but how?

These projects clearly work. Researchers need them, NSF understands this, and champions work to find mechanisms for continued funding.

- BLAS: 1979–present (47 years)
- CIG (Cyberinfrastructure for Geodynamics): ~25 years
- CSDMS (Community Surface Dynamics Modeling Systems): ~25 years
- REDCap (Vanderbilt): Consortium model combining academic & industry funding
- National survey & major data collection waves
- Telescopes & facilities

But when you ask why they succeeded, the answer is usually: "because the community cares" or "because the PI kept fighting for funding"—not a replicable framework.

Funding Maintenance

KEY RECOMMENDATIONS

Separate "maintenance of used infrastructure" from innovation funding. Fund it explicitly, on **contingent** multi-year timescales (20–30 years), based on demonstrated use. Analogy: observatories, long-term field measurements, HPC centers.

Systematize knowledge about long-term success. Repository of software success stories. Keynote talks. Workshop series: "how we built it and kept it going."

Mandate RSE career tracks as funding requirement for large CSSI awards. Universities respond to funding mandates. Explore hybrid teaching-maintenance roles. Consider NSF-hosted pool of expert RSEs available for temporary pairing.

Use NSF leverage to push communities toward shared standards (BRAIN Initiative model), sustainable models, good practices (CI/CD requirements are great!)

Maintain pipeline of funding for new & growing projects: Preserve diversity of models within domains. Resistance to consolidation; multiple approaches have value.

One Crazy Idea - “Software Center” Accounts

A common problem: A research project depends on a critical piece of software maintained by another team. They are **ready and willing** to direct $5-10K of their budget to supporting that software. The overhead to make this happen is astronomical:

- Subawards and university collaborations take months, and sometimes result in double overhead
- Founding an external organization (nonprofit or for-profit) is a huge undertaking

Allow NSF-funded research projects to allocate 5–10% of research grants toward maintenance of underlying software infrastructure, without this counting as indirect cost.

Aggregate these into NSF-administered "Software Infrastructure Center" accounts.

Just 5-10 projects contributing to a niche software package they rely on can provide the base of support for **maintenance** of key projects, distinct from the other important software priorities NSF funds

We don’t know *which* software needs to be maintained long term, but we know that *some* software will need it

“Software Centers” could:

- Decouple infrastructure health from individual PI discretion - PIs are used to “free” open source software, but this changes the question to **which** research software needs your support
- Lets the scientific research enterprise vote with their feet
- Prevents overhead bloat from eating maintenance budgets
- Creates multi-institutional buy-in

How to start? Which projects can have an account? Beyond direct input from funded projects, CSSI, PESOSE, etc. already do this!

Create dedicated sustainability mechanisms that work within existing systems

Thank you from the discussants!

Snapp-Childs, Winona
Kingsbury, Ryan
Sherrill, David
Dave May
Mark Perri
Daven Quinn
Mark Behn
Holden, Bradford
Caterina Strambio-De-Castillia
Michael Ryan
Hardy, David
Kline Struhl, Melissa
Pah, adam
Kobus Van der Merwe
Christian Escobar
Curtis McCully

Petar Maksimović
Xipeng Shen
Philip Ilten
Zehua Chen
Victoria Coles
Vincent Weaver
Bimal Gaudel
Jiannan Tian
Sandra Gesing
Joseph Tulenko
Gerard Lemson
Atanas Rountev
Satish Nair
Niaz Abdolrahim
Tianyuan Zhang
Xu Liang

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

Current challenges and concerns with AI and quantum computing

Forest Glen breakout

Facilitator: Andrew Rosen Scribes: Sean Cleveland, Emma Lejeune Reporting: Jaime Cernuda

Education

Trust and data security

Benchmarking

Quantum

Education: what the room is worried about

- LLMs let students bypass productive struggle; assign-and-return homework no longer teaches
- Students ask the AI instead of each other, so peer collaboration is eroding
- Verification is the missing skill: knowing when the AI is wrong and saying so
- Underneath it all: what is the goal of education now?

Education: where the group sees the opening

- Fund education research on AI and share the data openly so the community can learn what works
- Center critical thinking in the curriculum: pair programming, code review, harder and more realistic assignments
- Basic LLM literacy for every faculty member, with courses that adapt as evidence comes in

Action item

Make AI plus critical-thinking literacy a national requirement for college graduates

Measure of success: assessment of critical thinking in the context of AI

Education

Trust and data security

Benchmarking

Quantum

Trust and security: you can only trust what you understand

- We cannot say what a model will do or why; without that understanding, trust has nowhere to stand
- Sensitive data and prompts leave the institution for closed providers; who owns the outputs, and what if the provider disappears?
- A downloaded model cannot be told apart from a malicious one
- Agents need identity, credentials and scoped access, and the pieces of the harness are not designed together

Trust and security: where the group sees the opening

- Explainability and attribution as the core research problem: which prompt, retrieval or tool call drove a decision, and why
- Provenance and versioning of intermediate model state as the substrate for auditability
- Trust what you can inspect: open-weight and local models, signed model files, self-describing data capsules

Action item

Explainability and understanding as the source of trust and security

Measure of success: models and agents whose decisions can be traced to prompts, data and tool calls, and audited after the fact

Action item

Stop AI from acting human

Measure of success: open-weight models as capable as market models, with responses that explain rather than persuade

Education

Trust and data security

Benchmarking

Quantum

Benchmarking: what the room is worried about

- No one has agreed what error rate is acceptable for a given scientific task. 0%? 30%?
- Scientists are not building the benchmarks, and the ones that exist die with their paper
- Domains have not converged; many groups work in parallel with no shared standard
- Models learn the test: parameters leak and replication studies are poorly designed

Benchmarking: where the group sees the opening

- The CI community leads scientist-driven, domain-specific benchmarking of AI tools
- A benchmark is a maintained research product, hosted on national platforms with its datasets
- Borrow industry scoring and security evaluation frameworks; withhold test parameters from models

Action item

NSF and agencies fund domain-specific benchmark creation

Measure of success: maintained, well-documented, robust benchmarks that make it into the scientific literature

Education

Trust and data security

Benchmarking

Quantum

Quantum: what the room is worried about

- Quantum advantage is oversold; hardware is the hard limit and has been for a long time
- If the promises do not pan out, will agencies and the public back away from the field?
- Many in the room feel clueless: hardware is hard to reach, code is hard to compile, industry offers a service, not source
- Post-quantum encryption: how do we migrate existing systems and workflows?

Quantum: where the group sees the opening

- Find the problems where quantum is the right tool: computing quantum phenomena, not accelerating classical workloads
- Make hardware accessible and flatten the learning curve, with training for students and faculty alike
- Scientific benchmarks built by scientists, as with AI
- It is not too late to regulate, because the capability does not exist yet

Action item

Integrate with the DOE-led Genesis Mission, and reskill the cyber-workforce on AI and agents

Measure of success: joint calls and topic intersection; numbers trained and jobs created

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Future Directions for Cyberinfrastructure

Prasad Calyam, Daniel Aliaga, Md Tanvir Arafin

Breakout Session Summary, September 15th 2026

Cyberinfrastructure Definition is Evolving

- **CI scope is broadening:** HPC + AI workloads + edge computing + streaming data + hybrid campus/cloud + security standards
- **Services and users are changing:** beyond traditional client-server to cloud; increase in edge, mobile, and distributed environments
- **Rise of AI-enabled CI:** potential to boost productivity and creativity across domain sciences
- **Community CI:** greater investments for broader access, secure/scalable AI, inter-institutional data movement/sharing, workforce training

Current CI Community Status

- **Community has a strong foundation:** shared computing, science gateways, identity and access management, training programs, sustainable software, and curated domain science datasets
- **CI is more than hardware:** sustaining software, data, shared services, models has become a critical community need
- **Different users need different CI:** experts need advanced federated services; broader community/prosumers need usable, customizable tools
- **Emerging gap with Agentic AI:** today's human-oriented CI interfaces create friction for AI agents driving scientific workflows; alternately, AI agents can disrupt current CI interfaces

Desired Outcomes / Future Directions

- Prepare **NSF-supported CI for agentic AI** and **scientific workflows with human + AI teaming** – both from research support and workforce training perspectives
- **Standardize and validate AI-enabled workflows** before deployment; emphasis on trustworthy and interpretable approaches
- Strengthen **cybersecurity for next-generation CI**, especially edge and operational technology, defenses against AI-equipped threat actors
- Improve **collaboration and coordination of projects across NSF–DOE communities** - universities/labs, regional/national resources
- **Prepare for quantum computing and support diverse pent-up demands** such as inference services, local models, and agile datasets management

Recommendations for the Desired Future Outcomes

- **Launch OAC FutureCore program:** focused on foundational CI research for the age of agentic AI, emerging quantum computing
- **Support Agentic CI foundations:** support research in agentic: software development, data generation, science gateways, data management, and testing
- **Grow Edge CI capabilities:** establish standards/tools to provision and secure field science using e.g., drones, sensors, cameras, wireless networks, and tactical devices
- **Increase support for academia–industry centers of excellence:** accelerate building of CI for LLM inference, quantum computing
- **Foster more AI-ready datasets availability:** to enable linking and leveraging heterogeneous data across disciplines
- **Create Future CI Workforce:** fund training for AI-based scientific workflows across local and hyperscaler infrastructure

What Success Looks Like

- Scientific datasets become easier to connect, reuse, and combine across domains
- Domain scientists can deploy AI workflows easily on appropriate infrastructure
- Trustworthy agentic mechanisms become broadly applicable across domain CIs
- Secure edge CI becomes easier to design, deploy, and manage in the field
- NSF and DOE investments become more coordinated and complementary

---

- **Bottom line: build the foundational CI of the future that enables the next generation of AI-enabled scientific discovery**

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

Growing the Cyberinfrastructure Workforce

Educating students, training domain scientists, and sustaining the people who keep systems running

Nearly every discipline now depends on computational work, so gaps in this workforce become limits on the pace of research overall, and on who gets to participate in it.

THE CHALLENGE

One pipeline, three groups of people

Students

Educated in programs across the disciplines that now depend on computational work.

Domain scientists

Use CI daily without formal computing training.

CI professionals

Keep research computing running at centers that are already thinly staffed.

WHY NOW

AI is reshaping what all three need to know faster than curricula, hiring practices and career structures can adapt, while raising expectations for what one researcher can produce.

CURRENT STATUS AND DESIRED OUTCOMES

Where we stand, and what is still missing

WORKING TODAY

Proven venues

Campuses, summer schools, workshops, hackathons, internships.

Material already built

NAIRR, ACCESS, Jupyter, Binder, science gateways, all scattered across projects.

A lower barrier

Open OnDemand reaches HPC through a browser, not a batch scheduler.

The talent

Students arrive fluent with technology; staff often need only a short round of training.

Limits made visible

Instructors use LLMs on problem sets where answers are easy to verify, so students see where the models fail rather than being told about it.

STILL MISSING

- Real career structures for CI professionals, including hybrid roles split between running systems and doing domain science.
- Agreement on what CI foundations actually are.
- Training that reaches domain scientists, starting with a Python and data course for people who have never written code.

1

RECOMMENDATION 1 OF 3

Build shared AI infrastructure, and teach with it

WHAT TO DO

Stand up national secure AI models

Shared model infrastructure the whole community can teach and build on, rather than each campus improvising its own.

Ship classroom-ready tools and data

FAIR datasets with proper metadata, aimed at specific skills like anomaly detection and discipline-specific work.

Redesign courses, and support the faculty doing it

Rebuild courses around modern AI tools including agentic workflows, with continuing education so instructors can actually teach them.

HOW WE WOULD MEASURE IT

- Users and hours on the national AI models
- Datasets published, and how often each is adopted elsewhere
- Courses redesigned; uptake of agentic training activities
- Whether the educational tools hit their stated benchmarks

2

RECOMMENDATION 2 OF 3

Create real career pathways, and keep people in them

WHAT TO DO

Document how people actually get in

Map the real entry routes into the CI profession and what people run into along the way.

Build the missing on-ramps

Internships, apprenticeships and co-ops at a range of experience levels, plus capstone projects.

Make the jobs worth staying in

Flexible postdoc and RSE roles with at least 20% of time free of the funding grant; work with industry on what it needs; cut administrative red tape.

HOW WE WOULD MEASURE IT

- Longitudinal tracking through education, training, first jobs and moves
- Co-op and internship participation, and industry placement afterward
- Retention in flexible roles vs. inflexible ones; papers outside grant scope
- Industry satisfaction with hires, plus hiring statistics
- Staff time spent on non-value-added activity

3

RECOMMENDATION 3 OF 3

Make training stick, and judge it on outcomes

WHAT TO DO

Fund the workshop, then follow up

Keep funding in-person workshops and hackathons, but add structured follow-up, for example a one-day check-in three months later.

Solve the problem, don't hand it to an AI

Push trainees and new hires to solve problems themselves rather than handing them to AI.

Assess capability, not satisfaction

Judge programs by what participants can do afterward, not by end-of-course surveys.

HOW WE WOULD MEASURE IT

- Attendance and retention at the follow-up check-ins
- Open-source implementations of basic solutions contributed by trainees
- Capstone outcomes and national hackathon results as quality measures
- Standardized testing as a harder option in the US context

WHAT IT ADDS UP TO

Three moves, one pipeline

1. **Shared AI infrastructure and teaching material**

   National secure models, FAIR datasets, and courses rebuilt around the tools people actually use.

2. **Career pathways that people stay in**

   Documented on-ramps, internships and co-ops, and flexible roles with time outside the grant.

3. **Training judged on what people can do**

   Structured follow-up after workshops, real problem-solving, and assessment by capability.

Start here

A shared compendium of existing training modules from NAIRR, SCIPE, CSSI and similar programs, so nobody rebuilds what already exists.

2026 NSF CSSI / Cyber Training / SCIPE PI Meeting

Sept. 13 - 15, 2026 · Rockville, MD

Measuring and Communicating Scientific Impact: Challenges and Opportunities for NSF Broader Impacts

presented by: Martin Slawski

Background

- There is no universal definition of broader impact

- **Impact can be measured along a variety of axes**:
  - (trans-) disciplinary scientific impact
  - workforce development
  - curriculum and educational development
  - impact on public policy
  - national security and resilience
  - ….

- **Different stakeholders place different emphases**
  - Program directors: measurable outcomes and deliverables
  - Research communities: scientific advances
  - General public: societal relevance

Current Challenges in Measuring Impact

- Current reporting mechanisms tend to be confined to the duration of the funded project.

- Impact may not materialize during the funding period or immediately after. Conversely, impact may not be sustained.

- Conventional measures are based on easy-to-obtain numbers such as citations, downloads, conference attendance, etc.
  However, those can be noisy and do not adequately capture the full range of possible impacts.

Why Existing Metrics Are Insufficient

- Citations do not necessarily reflect meaningful use; AI-assisted writing further complicates the interpretation of citation counts.

- Similarly, #downloads may reveal little about adoption or the value that is created in the process.

- Downstream impacts are often invisible.

  Example:
  student uses software developed through a CSSI project
  → develops valuable skills → finds employment

- Long-term impacts are rarely tracked

Future Directions

- Understand who uses research products, how they are used, and what outcomes result; reconstruct downstream pathways of influence
- Conduct five-year or longer follow-up studies to determine whether projects remain active and impactful
- Track student and postdoc career trajectories (placement, productivity, leadership roles)

Recommendations

1. Improve software and infrastructure tracking (e.g., DOIs and stricter citation requirements for software).
2. Invest in impact-tracking cyberinfrastructure.
3. Track career outcomes of project participants (possible data sources: LinkedIn, Google Scholar, ORCID, …).
4. Complement metrics with qualitative evidence (e.g., via surveys).
5. Implement long-term follow-up assessments.
6. Expand evidence of societal impact (e.g., press releases, media coverage).
7. Develop guidance and standards for measuring and reporting broader impacts.