



August 24, 2026

Suzanne H. Plimpton
Reports Clearance Officer
National Science Foundation
Randolph Building, 401 Delany Street
Alexandria, Virginia 22314

Via: regulations.gov, Docket No. NSF-2026-OTR-0001
Information Collection 3145-0058

RE: Comments of the Computing Research Association on the Draft NSF Guidance on Financial Assistance, 91 FR 38030 (June 24, 2026), Docket No. NSF-2026-OTR-0001, Information Collection 3145-0058

I. Introduction, Organizational Interest, and the Scope of These Comments

The Computing Research Association (CRA) submits these comments in response to the National Science Foundation's request for public comment on the draft [NSF Guidance on Financial Assistance \(GFA\), published at 91 FR 38030 on June 24, 2026 under Information Collection 3145-0058](#). The draft would replace the NSF Proposal and Award Policies and Procedures Guide (PAPPG), version 24-1, with a new 203-page document. The PAPPG is the operating rulebook for nearly every federally funded academic computing research group in the United States. NSF has published the draft as a Paperwork Reduction Act information collection renewal and describes the revision as plain language, clarification, and the removal of content no longer applicable. It is considerably more than that.

Founded in 1972, CRA unites industry, academia, and government to advance computing research. CRA's membership includes more than 200 North American organizations active in computing research, including academic departments of computer science and computer engineering, research laboratories and centers in industry and government, and affiliated professional societies.

CRA makes two recommendations. The first recommendation is NSF should not finalize the GFA until the Office of Management and Budget (OMB) concludes its rulemaking on the [Guidance for Federal Financial Assistance, 2 CFR Part 200](#). NSF states in its own Summary of Changes that it drafted the GFA to align with OMB's proposed revisions to that guidance, and four individual guides repeat the point. Those revisions are not final. Second, and independent of how the OMB rulemaking concludes, we recommend that NSF withdraw or revise specific provisions of this draft that would damage US computing research.

NSF is the most important federal funder of fundamental computing research in the United States. The Directorate for Computer and Information Science and Engineering supports the majority of federally funded academic computing research, and NSF policy governs the daily operation of nearly every academic computing research group in the country. The research that policy governs has produced the internet, GPS, modern cryptography, and the foundational

advances in machine learning that underpin today's artificial intelligence industry. Damage to the system that produces that fundamental research would likely result in loss of leadership in the AI and cybersecurity space exactly when America needs it most.

[CRA filed comments](#) on July 13, 2026 in OMB's related rulemaking on the Guidance for Federal Financial Assistance, 2 CFR Part 200, Docket OMB-2026-0034. Those comments addressed many of the provisions this draft implements, and CRA incorporates them here by reference. Where this letter refers to CRA's position on a provision of the proposed OMB rule, that position is stated in full in the July 13 filing.

How these comments are organized

The draft GFA runs 203 pages and CRA has concerns about provisions in most of its 26 guides. This letter does not attempt to address all of them.

These comments are organized in two parts. Section II raises process concerns, meaning concerns about the way NSF has structured this proceeding: what NSF has told the public the draft does, what NSF has withheld, and whether comment on this record is possible at all. Those concerns come first because they are logically prior to the substance. If NSF has not disclosed what it is doing, no commenter can respond to it. Section III raises policy concerns about the substance of the draft. The three concerns in Section III A are ones CRA expects the bulk of the research community to share, and CRA states them briefly. The concerns in Section III B bear specifically on computing research, and CRA develops those at length, because they are the least likely to reach NSF from any other source and because several of them appear to affect the field in ways the draft's authors may not have anticipated.

Consistent with NSF's instruction, CRA begins each comment with the relevant guide and section reference in brackets. Page citations are to the draft's own printed page numbers.

II. Process: Whether Comment on This Draft Is Possible

The concerns in this section are not about what NSF has proposed. They are about whether the public can understand what NSF has proposed.

1. The draft is premature. NSF states in its Summary of Changes (draft pp. 5-6) that the GFA was drafted to align “with the proposed revisions to the Office of Management and Budget's (OMB) Uniform Guidance at 2 CFR 200 through the proposed rule published on May 29, 2026.” Guides 1, 3, 12, and 25 repeat the point. OMB's comment period closed on July 13, 2026 and no final rule has been issued. If OMB modifies or withdraws the provisions NSF has implemented, the GFA will be misaligned on the day it takes effect and institutions will absorb a second compliance rewrite. **CRA recommends that NSF hold the draft, conform it to the final rule, and re-open comment.**

2. A Paperwork Reduction Act notice is the wrong vehicle. NSF describes this revision as plain language, clarification, and the removal of content no longer applicable. However, the draft also alters merit review, creates a new authority to remove Principal Investigators, authorizes reduction of committed funding, and converts an institutional certification into a categorical eligibility bar. Those are substantive policy changes, and a burden-estimate notice is not an adequate process for them. **CRA recommends that NSF separate the substantive changes and propose them through notice and comment.**

3. The Summary of Changes does not adequately describe the draft. It omits significant changes, including the deletion of the Broader Impacts criterion's operative content and the new

national interest test applied to all foreign collaboration, and it describes Guide 8, which adds a new funding approval gate, as having “No significant changes.” It also announces five changes not found within the operative text: the TRUST research security framework, a publication cost disallowance, an artificial intelligence provision in research misconduct, and alignment with Executive Orders 14222 and 14332. CRA verified each by full-text search of the draft. **CRA recommends that NSF publish an accurate summary of changes, including a redline against PAPPG 24-1, and provide a sufficient notice and comment period once the accurate summary is available.**

4. Operative rules were not published with the draft. The draft makes at least seven documents controlling without reproducing them: (1) the Prior Approval Matrix, which Guide 12 designates as the exhaustive statement of what requires NSF's advance permission; (2) the full definition of a malign foreign talent recruitment program, deferred to NSF.gov and backed by an annual individual certification; (3) the NSF general conditions; (4) the indirect cost rate submission procedures, marked “(to be updated)”; (5) the TRUST framework; (6) the Gold Standard Science Implementation Plan; and (7) the enumerated societal outcomes that give Broader Impacts its content. Without access to these controlling documents, the public cannot provide relevant feedback and comment.

CRA adds one observation specific to its members. Four of those seven unavailable documents concern research security. The result is that **the draft does not disclose NSF's research security framework**, while that framework governs whether computing departments can recruit international graduate students and postdoctoral researchers and sustain international collaborations. Comment on an undisclosed framework is not possible. **CRA recommends that NSF publish every document the GFA makes operative and commit to notice and comment before revising any of them.**

5. The burden estimate is incomplete. NSF estimates 120 hours per proposal and 43,000 proposals annually. That measures proposal preparation and appears to exclude the recurring, post-award, and institution-level obligations the draft adds, among them annual malign foreign talent recruitment certifications, research security training, collection and institutional review of foreign employment agreements, institution-wide foreign gift aggregation with four-year retention, and development plan records for every funded student and postdoctoral researcher. Because the Prior Approval Matrix is unpublished, no commenter can estimate the burden of the prior approval regime at all. **CRA recommends that NSF publish a revised estimate, disaggregated by requirement, and re-open comment on it.**

6. NSF does not identify what text implements two of the Executive Orders it cites. The Summary of Changes states that the GFA aligns NSF policy with Executive Orders 14222, 14332, and 14303. The concern here is not the availability of those published orders, but rather that NSF has not said which provisions of this draft carry them out. Moreover, the draft's own practice shows an understanding of how to ground the draft in EOs. EO 14303 is cited in the operative text of Guide 13 and in that guide's Authorities list. EO 14292 is cited eight times across Guides 6 and 20, including in the screening procedures. Despite these examples, EO 14222 and EO 14332 appear once each, in a single Summary of Changes sentence, and in no Authorities list in the document. Nor is the implementation inferable: where EO 14247 is likewise cited only in the Summary of Changes, a reader can confirm what NSF did by observing that references to paper checks are gone. Nothing comparable is locatable for EO 14222 or EO 14332. This matters most for EO 14332, which concerns oversight of discretionary grantmaking, because the only provision that plausibly implements it is the unexplained approval gate discussed in Section III B. **CRA recommends that NSF identify, provision by provision,**

what text implements each Executive Order it cites, and re-open comment on those provisions.

III. Policy: The Substance of the Draft

Section III A states the substantive concerns shared broadly by the research community. Section III B develops the concerns specific to computing research. CRA presses both regardless of how NSF resolves Section II.

III A. Policy concerns the research community broadly shares

CRA states these three concerns briefly rather than at length. CRA expects them to be widely shared across the research community, and raises them here to record its own position.

1. Award stability and due process. A new “NSF Programmatic Stewardship Authority” (draft p. 125) permits NSF to remove a Principal Investigator, reduce funding, or terminate an award after consulting the institution’s administrative office but not the researcher, with “if necessary” as the entire trigger and no finding, notice, response, or appeal. NSF may reduce committed continuing grant increments “to respond to broad funding changes or new priorities” (p. 161). NSF may terminate immediately whenever “NSF believes” action is needed to protect the interests of the government, with the notice-and-cure step expressly removed (p. 198). Costs are unallowable while a recipient contests any of this, the review takes at least 105 days, and nothing stays the underlying action. Guide 26 then excludes specific conditions under 2 CFR 200.208 and every prior approval denial from review, and provides no appeal at all from a declination that is neither financial nor administrative (pp. 201-202).

CRA notes only the field-specific consequence. In computing, the Principal Investigator is frequently the only person holding the required expertise and supervising the trained doctoral students, so removing the investigator ends the project. Even if a recipient prevails after 105 days of unallowable costs, they will recover an award whose students have dispersed and whose reserved compute is gone. **CRA recommends written findings of cause, notice to the researcher as well as the institution, an opportunity to respond before action, preservation costs during a contested action, and a stay pending timely review.**

2. Retroactive expansion of award conditions. The order of precedence at Guide 10 (pp. 116-117) incorporates the NSF general conditions and the GFA by reference without specifying that the version in effect on the award date governs, and places award amendments first in precedence with no constraint on their content. NSF version-locks the cost principles at Guide 12 A.3 and nothing else. **CRA recommends extending that same lock to the general conditions, the GFA, and the Prior Approval Matrix.**

3. Merit review criteria and participation. The draft reduces the enumerated merit review questions from five to four, and the question removed is the only one that referenced societal benefit; retains Broader Impacts by name while deleting the enumerated societal outcomes that gave it content; removes “broadening the participation of groups underrepresented in STEM” from the Synergistic Activities examples while leaving the surrounding list intact; and reduces NSF’s nondiscrimination guidance to a single sentence while retaining the underlying statutory obligations and NSF’s authority to conduct compliance reviews. The terms “broadening,” “underrepresented,” and “diversity” appear nowhere in the 203-page draft. None of these deletions is disclosed in the Summary of Changes.

Two points CRA presses in its own voice. NSF continues to require an institutional certification of support for the Broader Impacts portion of every proposal under Section 526 of the America

COMPETES Reauthorization Act (draft p. 36), so NSF is requiring compliance with a statutory criterion whose content it has deleted. And NSF attributes the nondiscrimination removals to revoked Executive Orders, an explanation that does not reach Section 504 coordinators, Title IX coordinators, or age discrimination self-evaluations, which rest on statute and standing regulation. **CRA recommends that NSF restore the threshold review question and an enumerated, published list of qualifying societal outcomes, and identify the specific legal authority it believes required each nondiscrimination removal.**

III B. Policy concerns specific to computing research

[G8.A.Procedures 1(b) and G8.F] The “NSF leadership official” approval gate (draft pp. 104-105, 108)

CRA recommends that NSF define the term “NSF leadership official,” state the criteria that govern the approval, require a written statement of reasons whenever the official departs from the program officer’s recommendation, disclose that statement to the proposer, and report annually to the National Science Board on the number and disposition of such departures. CRA does not ask NSF to remove the step, only to make it accountable. This is CRA’s principal objection to the draft.

The draft provides:

“If the cognizant NSF Program Director recommends the proposal to move forward in processing and is approved by the NSF leadership official, it proceeds to the cognizant Grants Officer who generally make award/decline decisions within 30 days, subject to additional reviews for special cases.”

Guide 8 F describes the same architecture a second way, stating that funding recommendations “are reviewed by NSF leadership and then evaluated by NSF Grants Officers for financial, policy, and risk considerations,” and referring in its procedures to “the supervisory senior scientist” and to “Senior program staff.”

CRA wants to be precise about the scope of this objection. **CRA does not contend that agency leadership must accept every panel ranking, and does not ask NSF to give up programmatic discretion.** Program officers, division directors, and advisory bodies have always adjusted review outcomes for sound reasons, including balance across subfields, portfolio risk, and the avoidance of duplicative investment. NSF is an executive agency accountable for the funds it awards, and leadership involvement in discretionary award decisions is a normal feature of federal research management.

CRA’s objection is to three features of how the provision is drafted. First, the office is undefined: the phrase appears exactly once in 203 pages, no organizational level is specified, and no relevant scientific or programmatic expertise is required. Second, no criteria govern the decision: the approval is not tied to the merit review criteria in Guide 8 B, to panel rankings, or to any published standard, and where the draft does name a standard at a later stage it is “policy,” undefined. And third, no procedural discipline attaches: no timeline, no documentation, no statement of reasons, no notice to the proposer, and, under Guide 26 A, no appeal, because pre-award appeals reach only declines made for financial or administrative reasons by a Grants Officer.

The computing-specific consequence is anticipatory. Research on artificial intelligence safety, algorithmic transparency and auditing, content moderation systems, privacy-enhancing

technologies, and the societal effects of digital systems addresses subjects that are politically contested. Researchers reasonably calibrate proposals to what they expect will be funded, and faculty advise doctoral students on that basis. They should be able to expect that those decisions rest on merit and peer review. An unexplained approval gate over funding recommendations changes proposal behavior long before it is ever exercised, and that effect will appear in no statistics in any NSF reports.

CRA asks NSF to consider one further pairing. In the same document that creates this gate, NSF deletes protection from inappropriate influence from its definition of scientific integrity (draft p. 139), retaining the obligations that definition imposes on researchers while removing NSF's own commitment. Whatever NSF intends, the research community will read the two changes together. **CRA recommends that NSF restore the full scientific integrity definition and state expressly that NSF will protect scientific and technical judgments in merit review from inappropriate influence.**

[G12, with G4.B6 and G21.B] Immediate public access with no reconciled cost provision (draft pp. 7, 50, 130, 188)

CRA recommends that NSF delete the Summary of Changes statement that publication costs are disallowed, state affirmatively in Guide 12 that publication, open access, and research results sharing costs are allowable, and provide a mechanism for paying deposit and article processing charges for work accepted at or after the end of the period of performance.

CRA supports public access and has consistently urged that taxpayers be able to read the research their funds support. The objection is not to the mandate. It is that the draft states four different positions on whether the mandate can be paid for:

- Summary of Changes, p. 7: NSF “Disallows publication costs consistent with proposed revisions to 2 CFR 200.”
- Guide 12 Section E, pp. 134-135: no publication cost prohibition appears. The phrase “publication cost” occurs exactly once in the 203-page draft, in the Summary of Changes bullet above.
- Guide 12 A Procedures 2(c), p. 130: recipients may charge “allowable publication or research results sharing costs before closeout.”
- Guide 4 B6 Policies 3, p. 50: the budget may request funds including “Reports, reprints, page charges and illustrations.”

Guide 21 B adds a fifth, stating that third-party deposit fees “are still allowable expenses on NSF awards” (p. 188), in the same provision that removes the 12-month embargo and requires deposit at or before publication.

This matters more in computing than in most fields, because computing publishes primarily through peer-reviewed conference proceedings rather than journals, and the cost of publishing in those proceedings is not principally a publication charge. As of January 2026, ACM publishes all of its proceedings open access, and corresponding authors at the more than 2,700 institutions participating in ACM Open pay nothing, which covers roughly three quarters of papers at ACM-sponsored conferences. However, the cost did not disappear. Rather, it moved to a fixed annual institutional fee, with an article processing charge for authors whose institutions do not participate.

Moreover, computing proceedings require attendance at the relevant professional society conference. Acceptance obligates at least one author to register for and attend the conference, so the recurring per-paper cost of publication in this field is registration and travel, rather than traditional article processing charges (APCs). That is why CRA opposed the individual conference pre-approval requirement in proposed 2 CFR 200.432, and it is why NSF's decision not to adopt that requirement in this draft matters as much as it does. APCs do occur in computing research, but they arise principally in professional society journals, most of which now require open access publication, and those charges are squarely exposed to the disallowance the Summary of Changes announces.

[G4.B.Policies 2-3, G4.B6.Policies 2 and 5, with G6.F] No allowable cost path for cloud and accelerator computing (draft pp. 48, 50-51, 83-84)

CRA recommends that NSF add an express allowability path for externally procured cloud capacity, accelerator hours, and commercial model inference services; reverse the presumption against information technology equipment for computing research; and add CloudBank and the National Artificial Intelligence Research Resource to the list of national resources in Guide 6 F.

We also note a place where NSF can, while revising the PAPPG, address a known issue. CRA begins by acknowledging what NSF has built. CloudBank provides brokered commercial cloud access to eligible NSF researchers. NSF expanded it this year, and it works. CRA supports it. But CloudBank does not answer the question this guidance is supposed to answer. CloudBank credits are deliberately excluded from the formal proposal budget and are requested through a separate supplementary document, and eligibility is limited to solicitations that expressly provide for it. CloudBank is an allocation, not a cost rule. Nothing about it tells a sponsored projects office whether a Principal Investigator may place a commercial cloud invoice on Line G4 and charge it to the award.

That question is answered by Guide 4 B6 Policies 5, which permits computer service costs “only where it is institutional policy to charge such costs as direct charges” and requires “a justification based on the established computer service rates at the proposing organization.” That formulation presumes an internal campus recharge center with published rates. It contemplates neither a commercial vendor invoice nor a metered service.

CRA notes candidly that this language is carried forward from PAPPG 24-1. CRA is not asserting that NSF has made the situation worse. CRA's point is that NSF is rewriting this document specifically to reduce burden and remove content no longer applicable, and a cost provision drafted before commercial cloud computing existed is precisely such content. This is the moment to fix it.

The gap is not hypothetical. It reaches awards whose solicitations carry no CloudBank provision, capacity needed beyond a credit allocation, services CloudBank does not broker, and the multi-year reserved capacity a large training or systems project must plan against. Accelerator hours and inference capacity are now primary scientific inputs in artificial intelligence research, in the same sense that instrument time is in experimental physics, and the guidance does not say how to buy them.

Two related items. Guide 4 B Policies 2 states that “information technology equipment and systems are typically not eligible for direct cost support.” **We suggest that while IT equipment as part of administration might be “typically not eligible”, computing equipment for research, like other research equipment, should typically be eligible.** For a discipline

whose scientific instrument is information technology, a default against it inverts the ordinary allocability analysis and will generate justification burden on essentially every computing proposal. And Guide 6 F, a new consolidated section presenting the national resources available when local resources are insufficient, lists Frontera, ACCESS, NCAR resources, and PATH credits. It omits CloudBank, the program NSF built for this purpose, and the National Artificial Intelligence Research Resource. Adding two entries costs NSF nothing and would materially improve the guidance.

[G7.G.Policies 1 and G7.H.Procedures 1(e)] Conference and travel proposals (draft pp. 94-96)

CRA recommends that NSF revise the conference proposal eligibility provision to recognize that in some fields the regular meetings of professional societies are the primary peer-reviewed publication venue, and reframe the travel proposal code-of-conduct certification as a duty to verify and document rather than to certify a third party's policy.

The draft provides that conference proposals “will be supported only if equivalent results cannot be obtained by attendance at regular meetings of professional societies.”

That provision rests on an assumption that does not hold in computing. In most fields journals are the primary venue for peer-reviewed publication and society meetings are supplementary. In computing the relationship is reversed. The Conference on Neural Information Processing Systems, the International Conference on Machine Learning, the IEEE Symposium on Security and Privacy, ACM SIGCOMM, the ACM Symposium on Operating Systems Principles, and ACM CHI are some of the most consequential venues in their subfields, carry full peer review with acceptance rates far below many journals, and are meetings of professional societies. Read literally, the provision makes ACM and IEEE affiliated meetings the least fundable, inverting their scientific standing.

CRA recognizes that this language is carried forward from PAPPG 24-1 and is not a change NSF proposes here. CRA raises it because NSF is rewriting the document and because the provision has produced real confusion for computing proposers for years. A single clause acknowledging field variation would resolve it.

[G6.E.Policies 2 and G20.D.Policies 2] Dangerous Gain of Function provisions reach computational research (draft pp. 82, 182)

CRA recommends that NSF replace “is not prohibited” with “is not covered,” and state affirmatively that purely computational research, including the development, training, and evaluation of models, falls outside the definition of Dangerous Gain of Function Research and triggers no review, screening, or Institutional Review Entity referral. Oversight should attach at the point where computational output is used to synthesize or manipulate a physical agent.

CRA understands NSF to intend a carve-out for computational work. The text does not accomplish one. The draft provides:

“NSF will not fund research that would lead to dangerous gain of function. Proposals with potential dangerous gain of function research will be considered for funding only after NSF has conducted a review and assessed the associated risk. Purely computational (in silico) research that may include the development of computational models and software or may use such means to design novel forms of biological agents is not prohibited.”

Computational research is declared not prohibited. It is not declared outside the definition, not exempted from the pre-award review established by the sentence immediately preceding, and not excluded from the screening obligation in the accompanying procedures, which direct a Principal Investigator who is unsure to notify an Institutional Review Entity and contact the Program Director “to discuss whether the research can proceed.” A safe harbor that does not exempt work from the review is not a safe harbor.

The scope language compounds this. The definition's outcome list is keyed to what research “seek[s] or achieve[s],” and includes “facilitating their ability to evade detection methodologies” and “Generating or reconstituting an eradicated or extinct agent or toxin.” Both are readable as reaching computational work with no wet laboratory component.

The consequences are concrete. A group training a protein structure or sequence model, building artificial intelligence biodesign tools, or studying model capabilities in biology falls within this framework as drafted. Computer science departments would need access to an Institutional Review Entity, a biosafety-style committee most computing departments have no relationship with and that has no expertise in evaluating model architectures, training data, or evaluation protocols. A pre-award review of undefined duration is a substantial deterrent in a field where the frontier moves in months.

CRA supports oversight of biological risk and does not ask NSF to exempt work that produces physical agents. CRA asks NSF to place the boundary where the physical risk arises. The same policy language appears in Guide 6 E at draft page 82, so the correction must be made in both places; a change in only one guide would leave the provision unaltered in the guide proposers consult while preparing a proposal.

[G24.A.Policies 3] “Potentially classifiable” self-reporting with no fundamental research exclusion (draft pp. 195-196)

CRA recommends that NSF expressly restate the fundamental research exclusion of National Security Decision Directive 189 in Guide 24, and confirm that no NSF award creates dissemination restrictions on the results of fundamental research.

The draft states that NSF awards are for “unclassified, publicly releasable research,” and then requires that recipients “notify the cognizant NSF Program Director promptly if research results involve potentially classifiable information that warrants government restrictions on the dissemination of the result.”

CRA verified that the phrases “fundamental research,” “NSDD,” “controlled unclassified,” and “CUI” appear nowhere in Guides 22 through 26. NSF therefore affirms that its research is publicly releasable while placing on the recipient an affirmative duty to self-report when its own results are potentially classifiable, with no definition of that term, no standard for the judgment, no stated consequence, and no preservation of the exclusion that has governed this question since 1985.

This provision falls almost entirely on computing. Cryptography and cryptanalysis, offensive and defensive security research, side-channel analysis, autonomy, and frontier artificial intelligence capabilities research all produce results that someone could characterize as potentially classifiable. Asking a Principal Investigator to make that determination under threat of a compliance finding will not produce better classification decisions. It will cause researchers to avoid the areas where the judgment is hardest, which are the areas where an open research community is most valuable to national security.

[G14.A.Policies 3, with G14.A.Terms 3 and G2.C.Policies 4] Research security screening and the treatment of international collaboration (draft pp. 22, 141-142)

CRA recommends that NSF provide that any adverse research security determination will be disclosed in writing to the institution and the affected individual, will identify the information relied upon, will afford a documented opportunity to respond before action, will require human review of any automated indication, and will be subject to review. NSF should publish the criteria it will apply rather than reserving them. NSF should also confine the national interest justification to collaborations with entities designated under existing authorities.

On screening, the draft provides that “NSF reserves the right to perform risk assessments of proposals and awards using analytical tools to assess nondisclosures of required information from senior/key personnel, and establish other policies and procedures for identifying, communicating, and addressing security risks.”

CRA's concern is specific to how such tools behave on computing records. Screening of this kind operates on publication records, co-authorship relationships, affiliations, and funding acknowledgments. In computing, international co-authorship is the norm, papers routinely carry large author lists, author name disambiguation is genuinely difficult for common Chinese, Korean, and Indian surnames, and acknowledgment metadata is inconsistently structured across the conference venues where the field publishes. A system built on those inputs may flag researchers who have done nothing wrong, and crucially, this provision allows for the possibility that they will never learn about the false positive. The clause reserving authority to establish the substantive rules later means the criteria are not before the public at all.

On collaboration, Guide 2 C Policies 4 requires a justification that any foreign involvement is “in the national interest of the United States,” and asks proposers to assess the “administrative capacity” of the foreign site against a domestic benchmark. Proposed 2 CFR 200.220 confines its national interest test to covered foreign countries and entities. NSF has applied that standard to all foreign involvement of any kind, so a subaward to a Canadian university or a consulting agreement with a researcher in Germany now requires the justification. Guide 6 G adds a requirement to justify activities country by country with no standard stated, and Guide 3 H defines an international activity to include collaboration conducted from the United States “using virtual technologies” and travel to attend conferences, which describes ordinary computing practice.

Approximately 57 percent of U.S. doctoral degrees in computer science are awarded to temporary visa holders. Computing research capacity in the United States depends on the ability to recruit and retain those students and to sustain collaborations with allied-nation institutions. CRA supports research security measures directed at identified risks. However, measures that treat all international collaboration as presumptively suspect impose real costs on U.S. capability and produce no commensurate security benefit.

CRA supports and asks NSF to retain verbatim the three exclusions at Guide 14 A Policies 6, which preserve scholarly presentation and publication, participation in international conferences and reciprocal exchanges, and the advising of and writing of recommendations for foreign students. Those exclusions are what keep ordinary practice in the field lawful.

[G25.B.Policies 1] Research misconduct definition, and the absent artificial intelligence provision (draft p. 199)

CRA recommends that NSF restore verbatim the honest error exclusion and the “significant departure from accepted practices of the relevant research community” standard, state expressly that a failure to reproduce a result is not itself falsification, and publish for comment any provision concerning artificial intelligence tools before adopting it.

The draft restates a truncated definition of fabrication, falsification, and plagiarism while omitting the elements that make the standard administrable: the exclusion for honest error and differences of opinion, the requirement of a significant departure from accepted practices of the relevant research community, the intent element, and the preponderance of the evidence standard. Each appears in 45 CFR 689, which the guide cites, so this is a conformity request rather than a policy change.

The omission matters acutely in computing. Nondeterministic training runs, seed-dependent results, hardware-dependent performance measurements, benchmark and dataset contamination, and impacts of different accelerators mean that differing results are ordinary features of the field, and can be important insights in their own right. A definition that drops the honest error and accepted-practices limitations invites allegations founded on them, against exactly the large-scale machine learning work that is hardest to replicate independently.

On artificial intelligence, NSF's Summary of Changes states that Guide 25 “Clarifies research misconduct requirements by including reference to artificial intelligence tools.” CRA verified that no reference to artificial intelligence, machine learning, or generative tools appears anywhere in Guide 25 or in Guides 22 through 26. CRA cannot comment on unpublished text, but can state the requirement in advance. Any such provision must distinguish three things: using a language model in the drafting and editing process, using artificial intelligence as a research instrument, and studying artificial intelligence systems, where model outputs are the primary research data. For CRA's members the second and third are the discipline itself. A provision that fails to separate them would render a large share of computing research presumptively suspect, and NSF should not adopt one without publishing it for comment.

[G13.D] Gold Standard Science as a proposal evaluation standard (draft p. 140)

CRA recommends that NSF either retain Gold Standard Science as a statement of agency practice and remove it as a proposal evaluation standard, or, if NSF intends to apply it to proposals, state who applies it, publish field-specific interpretive guidance developed with the research community, state the consequence of an adverse determination, and provide a route to contest one.

Guide 13 D states nine tenets and one sentence of policy: NSF “expects the highest standards of scientific rigor, integrity and adherence to tenets of Gold Standard Science in proposals, as appropriate for the field of science and research modality.” Guide 13's introduction directs Principal Investigators to use the guide to understand NSF's expectations regarding “the practice of Gold Standard Science.”

CRA supports rigor, transparency, honest reporting of uncertainty, and freedom from conflicts of interest, and has no objection to NSF stating those values. The problem is the qualifier. “As appropriate for the field of science and research modality” is itself undefined, and the draft does not say who determines what is appropriate, at what point, against what standard, or with what consequence. An undefined qualifier applied by an unnamed evaluator is not a safeguard; it is unbounded discretion over a nine-part standard that is assessed “in proposals” and that does not appear among the merit review criteria in Guide 8 B.

The tenet illustrating the problem most clearly is “structured for falsifiability of hypotheses.” Theoretical computer science, algorithm design and analysis, formal verification, programming language semantics, and cryptographic proof construction do not proceed by falsifiable empirical hypothesis. They produce constructions, proofs, and artifacts whose correctness is established by proof rather than experiment. A reviewer applying that tenet literally would find a substantial share of computing research, including work that led to Turing Awards, to be non-compliant. The tenet “collaborative and interdisciplinary,” stated as an expectation, similarly disadvantages single-investigator theoretical work. CRA does not suggest NSF intends either outcome, but nothing in the text prevents it.

[G5.H and G5.C.Policies 6] Data management plans and the disclosure of donated computing (draft pp. 64, 71-73)

CRA recommends that NSF raise or remove the four-category limit in the Data Management and Sharing Plan webform, confirm that access limitations for security research data, human subjects data, and third-party proprietary data will be accepted without a heightened burden on the Principal Investigator, and either raise the in-kind disclosure threshold for computational resources or publish valuation guidance for them.

The draft directs proposers to “Add a maximum of four data or research product categories.” A single computing award routinely produces source code, trained model weights, benchmark datasets, hardware designs or artifacts, and instructional materials, each with different repositories, licenses, and retention profiles. Four categories will force proposers to omit products or to aggregate unlike things.

The access limitation question is more serious. Security research generates vulnerability and exploit data whose immediate release would cause harm, human-computer interaction research generates human subjects data that cannot be archived publicly, and industry-provided datasets carry contractual restrictions. The draft permits a proposer to state limitations but assigns the justification to a named individual under an “Accountability” field. NSF should confirm that responsible non-disclosure in these categories is recognized rather than treated as a deficiency.

Separately, the draft requires disclosure of in-kind contributions valued at \$5,000 or more, and defines in-kind contributions to include “equipment; data or data sets; supplies; other expendable property; goods and services.” In computing, donated cloud credits, accelerator cluster access, model inference credits, and proprietary datasets routinely exceed that threshold. The draft requires that each must be assigned a dollar value and a person-month commitment, but supplies no valuation methodology. Two institutions will inevitably value the same credit grant very differently, and under Guide 3 D, these good-faith valuations would be a disclosure violation if NSF later disagrees with the institution.

[G11.C.Policies 1] Changes to computing resources treated as changes in scope (draft p. 122)

CRA recommends that NSF limit this provision to changes that materially reduce the resources necessary to perform the approved scope of work.

The draft provides that “Changes to the Facilities, Equipment and Other Resources section of the approved award constitute a change in scope,” and changes in scope require prior NSF approval.

In computing proposals that section describes compute clusters, accelerator allocations, cloud credits, testbeds, instrument access, and data resources. Those change constantly, and usually for reasons outside the Principal Investigator's control: a cluster reaches end of life, an allocation is reassigned, a national facility revises its queue policy, a vendor changes terms, a collaborating institution's testbed goes offline. Under this provision each such event becomes a change in scope requiring prior approval even where the research plan is unaffected, generating approval traffic that consumes Principal Investigator and program officer time without improving stewardship.

IV. Provisions CRA Supports

CRA supports and asks NSF to retain: the open communication policy at Guide 23 A Policies 1; the three exclusions from the malign foreign talent recruitment program definition at Guide 14 A Policies 6; the prohibition at Guide 4 B9 Policies 1 on NSF staff suggesting that Principal Investigators reduce or waive indirect costs, which counters the indirect cost preference in proposed 2 CFR 200.205; the removal of research security concern as a ground for returning a proposal without review at Guide 9 B; proposer access to review materials at Guide 8 G; the recipient-executed first no-cost extension at Guide 10 B; the deferral of termination reporting until the recipient has exhausted review at Guide 25 A Policies 6; the increased thresholds for Planning, RAPID, and EAGER proposals at Guide 7; and the childcare and family care resource descriptions in conference proposals at Guide 7 G.

V. Conclusion and Recommendations

CRA appreciates NSF's effort to make its financial assistance policies clearer, and supports the reorganization the draft undertakes. But the draft implements a rule that is not final, makes substantive policy changes through a Paperwork Reduction Act notice, describes those changes inaccurately in its own summary, and leaves the documents that carry its research security framework unpublished. Several of its provisions were drafted without reference to how computing research is conducted, and would impose costs the draft does not appear to anticipate.

CRA recommends that NSF:

- Hold the draft GFA until OMB issues a final rule at 2 CFR Part 200, conform the GFA to it, publish every document the GFA makes operative, and re-open comment;
- Identify, provision by provision, what text implements each Executive Order the Summary of Changes cites, and re-open comment on those provisions;
- Define "NSF leadership official," state the criteria governing the approval, require and disclose a written statement of reasons for any departure from the program officer's recommendation, report annually to the National Science Board, and restore the full scientific integrity definition including protection from inappropriate influence [G8.A, G13.C];
- Delete the Summary of Changes statement that publication costs are disallowed, affirm that publication and open access costs are allowable, and provide a mechanism for work accepted at or after the period of performance [G12, G21.B];
- Add an allowability path for commercial cloud, accelerator, and model inference costs; reverse the presumption against information technology equipment; and add CloudBank and the National Artificial Intelligence Research Resource to Guide 6 F [G4.B, G6.F];

- Recognize in the conference proposal provision that professional society meetings are the primary peer-reviewed venue in some fields, and reframe the travel certification as a duty to verify [G7.G, G7.H];
- State in both Guide 6 E and Guide 20 D that purely computational research is not covered by the Dangerous Gain of Function framework, and restate the National Security Decision Directive 189 fundamental research exclusion [G6.E, G20.D, G24.A];
- Provide notice, correction, human review, and a route to contest any adverse research security determination and publish the criteria; and confine the national interest justification to entities designated under existing authorities, excluding remote collaboration and conference travel [G14.A, G2.C, G3.H, G6.G];
- Restore the honest error and accepted-practices limitations to the research misconduct definition and publish any artificial intelligence provision for comment before adopting it; and either define who applies Gold Standard Science with published field-specific guidance or retain it as a statement of agency practice [G25.B, G13.D];
- Relax the Data Management and Sharing Plan category limit, recognize responsible non-disclosure for security and human subjects data, publish valuation guidance for donated computational resources, and limit Facilities section changes treated as scope changes to those that materially reduce the resources needed [G5.H, G5.C, G11.C]; and
- Adopt the procedural protections on award stability, termination, and review described in Section III A.

Computing research is among the most consequential long-term investments the United States makes in its economic strength and national security. The federal research funding system, built over decades on peer review, scientific independence, and stable award conditions, produced the internet, GPS, modern cryptography, and the foundational advances in machine learning on which today's artificial intelligence industry rests. NSF built much of that system, and NSF policy has been its most important safeguard. CRA urges NSF not to weaken those safeguards in a document presented as a plain-language rewrite.

CRA and its member organizations appreciate NSF's consideration of these comments and stand ready to work with NSF on approaches to clarity and administrative simplification that strengthen rather than undermine the US research enterprise.

Respectfully submitted,

Computing Research Association
1828 L Street, NW, Suite 800
Washington, DC 20036
cra.org

Points of Contact:

Peter Harsha
COO & Senior Director of Government Affairs
harsha@cra.org

Brian Mosley
Associate Director of Government Affairs
bmosley@cra.org