RISC

Research Infrastructure Security Community

Comments on Section 5.3 of the 2025 Draft of the Research Infrastructure Guide (RIG)

6 March 2025

Craig Jackson, Deputy Director, Indiana University Center for Applied Cybersecurity Research; Senior Personnel, Trusted CI

Jim Berhalter, Director of IT, National High Magnetic Field Laboratory

Wade Craig, Information Security Officer, National Radio Astronomy Observatory

Brian J. Koopman, Associate Research Scientist, Advanced Simons Observatory

Sean Peisert, Senior Scientist, Berkeley Lab; Adjunct Professor, UC Davis; Director and PI, Trusted CI

Scott Russell, Chief Security Analyst, Indiana University Center for Applied Cybersecurity Research; Deputy Director, Trusted CI

Jerry Brower, Cybersecurity Team Lead, NSF NOIRLab

Eric Cross, Head of Information Technology, National Solar Observatory

Jelena Mirkovic, Research Associate Professor, University of Southern California

Craig Risien, CI Systems Project Manager, Ocean Observatories Initiative

Randy J Trudeau, Chief Information Security Officer, Laser Interferometer Gravitational-Wave Observatory

David Balenson, Interim Director, Networking and Cybersecurity Division, USC Information Sciences Institute

Sam Chan, IT Director and Chief Information Security Officer, Giant Magellan Telescope

Ken Feldman, Systems Lead, UW Oceanography; ARF SatNAG; OOI RCA

Christopher Morrison, Head of IT Operations, NSF NOIRLab

Kelli Shute, Chief Project Manager, Indiana University Center for Applied Cybersecurity Research; Executive Director, Trusted CI

Table of Contents

About RISC

The Research Infrastructure Security Community (RISC) is a community of practice facilitated by Trusted CI for NSF-community technology and security practitioners. RISC is composed of graduates of the Trusted CI Framework Cohorts,[1] who continue to gather as a community to expand their cybersecurity knowledge, share experiences, and build relationships.

RISC is composed of technology leaders and professionals who, in many cases, will be the ones tasked with implementing the cybersecurity requirements that appear in the NSF Research Infrastructure Guide.

Authorship

These comments were developed in consultation with the full membership of the Research Infrastructure Security Community. However, only those named individuals on the cover page should be considered commenters.

The views and opinions expressed herein are those of the individuals identified and do not necessarily represent the views and opinions of any entities they represent.

Acknowledgements

The Research Infrastructure Security Community was initiated and is convened by Trusted CI.

Trusted CI, the NSF Cybersecurity Center of Excellence is supported by the National Science Foundation under Interagency Agreement #A2407-049-089-064206.0. The views expressed do not necessarily reflect the views of the National Science Foundation or any other organization.

[1] For more on Trusted CI and the Trusted CI Framework Cohorts, *see* https://www.trustedci.org/.

<u>Comment #1</u>: NSF should continue to emphasize taking a programmatic, risk-based approach to cybersecurity.

The Research Infrastructure Security Community (RISC) strongly supports NSF's goals of advancing the cybersecurity and resilience of research infrastructure (RI) organizations in the newest draft Research Infrastructure Guide (RIG). As a community of cybersecurity practitioners at RI organizations, we are advocates for improving cybersecurity and we want to be a resource to NSF as it crafts and refines its approach to advancing the cybersecurity of the NSF RI ecosystem.

Draft Section 5.3 includes many elements we want to highlight and reinforce, as they align with and canonize practices the community has adopted organically. In particular, many of these practices can be seen in the Trusted CI Framework's "Musts," which we have highlighted below. The Trusted CI Framework is a voluntary framework for organizations to establish and refine their cybersecurity programs. The Trusted CI Framework is a product of Trusted CI, the NSF Cybersecurity Center of Excellence. To date, the Framework has been adopted by no fewer than 23 NSF Major Facilities and Mid-Scale Research Infrastructure projects, along with a growing number of other RI organizations.

In particular, we are encouraged to see that the draft RIG emphasizes:

- Taking a programmatic approach toward cybersecurity, including explicit use of the Trusted CI Framework's four Pillars: Mission Alignment, Governance, Resources, and Controls in Section 5.3.7.[2]
- Having leadership aware of and appropriately involved in cybersecurity decision making (5.3.1). This aligns with Must 5 (Leadership Involvement).
- Including cyber risks in the organization's risk register as warranted (5.3.4). This aligns to Must 6 (Risk Acceptance) and Must 11 (Adequate Resources).
- Creating a cybersecurity lead role with access to organizational leadership (5.3.5, footnote 3). This aligns with Must 7 (Lead Role).
- Having cybersecurity support the research mission (5.3.1), and empowering RI organizations to select the framework (5.3.7) and baseline control set (5.3.7.4) that makes the most sense for their mission. This aligns to Must 1 (Mission Focus) and Must 15 (Baseline Control Set).
- Acknowledging that "cyber risk is rarely eliminated …, and residual risk should be acknowledged and formally accepted by facility leadership." This aligns to Must 6 (Risk Acceptance) and Must 11 (Adequate Resources).
- Having a discrete budget for cybersecurity (5.3.3). This aligns to Must 12 (Budget).

Recommendations:

a) NSF should continue to emphasize taking a programmatic, risk-based approach to cybersecurity.
b) NSF should retain language that reinforces practices the community has rallied around organically, such as the Trusted CI Framework.

[2] The citation and DOI for the Trusted CI Framework Implementation Guide is: "Jackson, C., Cowles, B., Russell, S., Adams, E. K., Kiser, R., Ricks, R., & Shankar, A. (2021). The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators (1.0). Zenodo. https://doi.org/10.5281/zenodo.4562447".

<u>Comment #2</u>: NSF should return to its approach of identifying high-level cybersecurity guidance and concerns in the RIG, and increase its use of other mechanisms for discussing more detailed cybersecurity recommendations.

Draft Section 5.3 represents a substantial shift in NSF's approach to guidance on cybersecurity, and is likely to result in increased cost and confusion among RI organizations. The 2025 draft RIG greatly revises and expands "Section 5.3: Information Assurance" (previously titled "Guidelines for Cybersecurity of NSF's Major Facilities"). While draft Section 5.3 retains a good deal of sound high-level guidance from the previous RIG, it also introduces a number of additions and changes.

We want to emphasize that NSF's historic approach of focusing on high-level guidance, as opposed to prescribing specific requirements, has been successful for RI organizations. NSF facilities vary considerably in terms of their size, budget, mission, and risks, and benefit from greater flexibility when it comes to implementing their cybersecurity programs. While there are some fundamentals that are always true, and which are valuable to highlight and emphasize, overly prescriptive approaches unnecessarily burden these organizations.

Additionally, NSF has other mechanisms for providing guidance that are more flexible than the RIG, a document which is updated roughly once every five years. These include the NSF Research Infrastructure Workshop, NSF Cybersecurity Summit, direct engagement through program officers, and participation at RISC Workshops. We encourage NSF to build upon these to communicate evolving NSF interests, concerns, and recommendations. and to provide for a greater degree of dialogue.

Recommendations:

a) NSF should return to its historic approach of using the RIG to identify high-level guidance and concerns, and leave ***how*** to tackle those to the RI organizations.
b) NSF should use other mechanisms to engage with RI organizations about the detailed aspects of building and evolving their cybersecurity programs.
c) NSF should establish a periodic meeting between its internal cybersecurity committee and the cybersecurity leads of the Major Facilities and Mid-Scale Research Infrastructure.

<u>Comment #3</u>: NSF should avoid the term "Information Assurance" and return to using "cybersecurity" to set the scope of Section 5.3.

NSF uses the term "information assurance" to characterize the guidance in Section 5.3, stating "Information Assurance (IA) is used as the umbrella term inclusive of cybersecurity, data protections (including privacy), cyber risk management, and resilience." Although we appreciate NSF's goal of emphasizing the full scope of cybersecurity, we think that this change in language is likely to be more confusing than helpful. We recommend sticking with the more commonly used term "cybersecurity," and emphasizing that this term actually encompasses the broader scope reflected in the section.

Although seemingly just a change in terminology, switching from "cybersecurity" to "information assurance" will impose notable costs on RI organizations. These costs include:

- Effort to understand NSF's new definition and apply it to the scope of existing programs.

- Effort to retrain an entire workforce on this apparent change of scope. Many RI organizations have spent years maturing and socializing their cybersecurity programs, and making sure those programs reach the full scope of the organization's technology environment and workforce.
- Effort to update internal policies, procedures, and plans.

We also think the switch is likely to cause confusion for a few reasons:

- The definition of "Information Assurance" in this draft is **different from federally recognized definitions of the term**.[3]
- For at least the last 10 years, the larger cybersecurity community has been moving away from the term "information assurance" in favor of "cybersecurity." See, *e.g.*, DoD Instruction 8500.0, specifically 1.d.[4] For many, "information assurance" calls back to a time before the security risks surrounding information technology and operational technology were so pervasive. For others it connotes a focus on strict compliance as opposed to risk management. If NSF resurrects this old term, it poses the risk of confusion or misunderstanding.

Ultimately we feel that the term "cybersecurity," both as defined by the federal government[5] and as used in the active December 2021 version of the RIG, adequately reflects NSF's intended scope of "cybersecurity, data protections (including privacy), cyber risk management, and resilience."

Recommendations:

a) NSF should remove references to "Information Assurance" and return to using "cybersecurity."
b) NSF should use recognized federal definitions of key terms. Authoritative sources of relevant federal definitions can be found in NIST's Computer Security Resource Center Glossary and CNSSI 4009.

Comment #4: NSF should align the language in Section 5.3 with the usage laid out in Section 1.1, removing the words "obligation," "expectation," "requirement," "essential," and other similar words.

Section 1.1 of the draft RIG states:

> "The terms *must* and *should* are used consistently throughout this Guide, adhering to federal plain language principles. ***Must*** conveys an obligatory action or legal requirement by the Proposer or Awardee, whereas *should* signifies a strong recommendation or good practice, but not a mandatory requirement." [bold emphasis in original]

However, draft Section 5.3 uses additional terms that make it unclear what is required versus recommended. In addition to using the terms ***must*** and ***should***, the section also uses the following

[3] *See* https://csrc.nist.gov/glossary/term/information_assurance, accessed 24 February 2025.
[4] "Adopts the term "cybersecurity" as it is defined in National Security Presidential Directive-54/Homeland Security Presidential Directive-23 (Reference (m)) to be used throughout DoD instead of the term "information assurance (IA)."" at https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodi/850001_2014.pdf, accessed 6 February 2025.
[5] *See* https://csrc.nist.gov/glossary/term/cybersecurity, accessed 24 February 2025.

terms: ***obligations*** (*e.g.*, "Section 5.3.3 Awardee Obligations"); ***expects*** (*e.g.*, Section 5.3.6 "NSF expects these controls to be an integral part of the Awardee's management practices for each Major Facility and Mid-scale RI and to be targeted for prioritized implementation."); ***requires*** (*e.g.*, Section 5.3.1 "In pursuit of this goal, NSF requires sufficient documentation and evidence that Awardees are performing good stewardship regarding IA."); and ***essential*** (*see* Section 5.3.7.4 "In response to an analysis of historical incidents, NSF has identified a small set of essential security controls for prioritized implementation.")

Notably, the terms "obligation"[6] and "requirement"[7] are both defined in the draft RIG's Lexicon (Section 8.2: Terms and Definitions), but it is unclear if their usage in this section is based on those definitions.

This inconsistency and accumulation of terms is confusing regarding what parts of Section 5.3 are actually required. We are concerned this will result in unnecessary costs as RI organizations try to understand and align to this updated section. Additionally, other stakeholders (*e.g.*, program officers) will also be tasked with interpreting these terms, and may very well come to different interpretations than the RI organizations.

Recommendations:

a) NSF should align the language in Section 5.3 with the usage laid out in Section 1.1, removing the words "obligation," "expectation," "requirement," "essential," and other similar words.
b) NSF should consolidate the ***musts*** into a single subsection at the beginning of the draft Section 5.3, and clearly state that the remainder of the section is guidance on possible ways to implement those ***musts***.

<u>Comment #5</u>: NSF should remove the "NSF Critical Controls" and refocus its guidance on adopting and evaluating against a baseline control set.

In the draft RIG Section 5.3.6, NSF states:

> "Several cybersecurity controls have been identified as exceptionally impactful in improving the resilience of Major Facility and Mid-scale RI, as illustrated in Table 5.3.6-1. NSF expects these controls to be an integral part of the Awardee's management practices for each Major Facility and Mid-scale RI and to be targeted for prioritized implementation. These controls should be integrated into the design and construction of new RI, and phased into operational RI within a reasonable timeframe, as determined by the NSF Integrated Project Team (IPT)."

We think that introducing this new control set is unnecessary for the following reasons:

1) In Section 5.3.5, NSF encourages the adoption and use of a third party baseline control set.

[6] "Obligation" is defined as "[a] definite commitment that creates a legal liability of the government for the payment of goods and services ordered or received, or a legal duty on the part of the United States that could mature into a legal liability by virtue of actions on the part of the other party beyond the control of the United States." We interpret this definition to apply exclusively to the Federal Government, and therefore inappropriate to use when referring to Major Facilities or Mid-scale Research Infrastructure.

[7] "Requirement" is defined as "[a] condition or capability that is required to be present in a product, service, or result to satisfy a contract or other formally imposed specification."

We strongly agree with this practice, which is formalized as Must 15 in the Trusted CI Framework. Carefully selecting and thoughtfully using a baseline control set -- particularly one like the CIS Controls, which is itself already prioritized -- is more than sufficient to help RI organizations prioritize implementation in light of their distinctive missions and realities.

2) The language used to define and describe the 14 "NSF Critical Controls" does not quote or map directly to existing baseline control sets. In Table 5.3.6-1, while NSF references controls from two of NIST's baseline control sets, the language NSF uses differs from that found in the referenced federal source.
3) While NSF states that these 14 "NSF Critical Controls" have been "identified as exceptionally impactful," it describes no research or systematic study measuring impact. When NSF personnel have spoken publicly about these controls, they have described these 14 as practices that one or more NSF facilities have put into place ***after*** a security incident. If, in fact, that is the basis for their inclusion, then the claim that these are "exceptionally impactful" is misleading. Emphasizing controls implemented only ***after*** an incident **would necessarily exclude any highly impactful practices** already in place at those RI organizations prior to the incident. Likewise, the good practices in place at RI organizations that have not experienced an incident would also be excluded.
4) There exists a growing body of rigorous research into cybersecurity control effectiveness.[8] If NSF is going to recommend specific controls, we recommend relying on this body of research instead.
5) We believe the RIG is updated too infrequently for it to be an effective venue to emphasize prioritization of specific, detailed cybersecurity controls.

Recommendation:

a) NSF should remove the "NSF Critical Controls" and refocus its guidance on adopting and evaluating against a baseline control set.

<u>Comment #6</u>: NSF should clarify or remove references to "phishing-resistant" multifactor authentication.

Table 5.3.6.1 details the "NSF Critical Controls Set." Controls "NSF1" and "NSF2" both refer to "phishing resistant MFA." However, no definition is given for this phrase, and the references provided do not discuss "phishing resistant" MFA. We recommend removing this term, as we are unsure what it means in practice, and are concerned that implementing it may prove costly and impractical.

When the RISC group discussed this term, we were unable to come to a consensus on what we thought it meant. Some RISC members suggested that the term might simply mean avoiding MFA that relies on SMS, as this has been shown to be vulnerable to some attack types. Yet others identified guidance from the Cybersecurity & Infrastructure Security Agency (CISA) on phishing resistant MFA.[9] Notably, the CISA guidance claims that the only effective phishing resistant MFA is either "FIDO/WebAuthn Authentication" or "PKI-based MFA" (*e.g.*, physical MFA tokens or a card similar to the CAC used in the Department of Defense). Multiple RISC members noted that implementing this kind of authentication would be significantly more costly than widely accepted

[8] *See, e.g.*, discussion of the research that uncovered the "Transformative Twelve" controls at https://www.lawfaremedia.org/article/the-difficulties-of-defining-secure-by-design. Accessed 24 February 2025.
[9] *See* https://www.cisa.gov/sites/default/files/publications/fact-sheet-implementing-phishing-resistant-mfa-508c.pdf.

MFA solutions, and in some cases impracticable, considering that their science missions rely on users from other facilities distributed around the globe obtaining remote access to their resources.

Recommendations:

a) NSF should clarify or remove references to "phishing-resistant" multifactor authentication.
b) NSF should not include a recommendation or requirement for physical MFA tokens or cards to qualify as phishing resistant. These come with significant costs and impacts to mission, and there is limited evidence of their added value relative to other multifactor implementations (*e.g.*, authentication apps on smartphones).

We want to be clear that MFA is empirically-proven as one of the most impactful types of cybersecurity controls,[10] and we support emphasizing its importance. Similarly, we recognize that phishing attacks are a major vector for many cybersecurity attacks, including ones that seek to overcome MFA. However, consistent with other guidance in Section 5.3, RI organizations should take a risk management approach to cybersecurity. That means they need the flexibility to address risks in the manner that makes the most sense for their mission. Requiring specific implementations of MFA, particularly costly ones that are not commonly used in the community, denies RI organizations that flexibility. We encourage the NSF to emphasize the importance of MFA without prescribing specific solutions.

Comment #7: NSF should clarify that existing documentation can be used to satisfy the requirements in Section 5.3.

Draft Section 5.3 represents a substantial increase in the amount of new documentation required or recommended of Major Facilities and Mid-scale Research Infrastructure. While we understand that documentation plays a valuable role in cybersecurity operations, we think that NSF's needs can be adequately met by existing documentation commonly used by the community.

Much of this new documentation would come in the form of the new "Information Assurance Management Plan" (IAMP), described in Section 5.3.5. The IAMP has several component parts, including: 1) Risk Treatment Plan(s)[11]; 2) a description of "IA Program Operations," which includes a) "Programmatic Processes", b) "Baseline Security Functions," and c) "Supplemental Responsibilities of the Security Program"; and 3) an "Assessment Plan." Each of these components includes limited detail regarding implementation.

In total this amounts to a considerable increase in the documentation requested by the NSF, and will require considerable effort to generate and maintain. Nearly every RISC member has noted that effort limitations are one of the biggest challenges facing their cybersecurity program. These additional documentation requirements will materially impact RI organizations' ability to make other important changes to their cybersecurity programs.

Additionally, it is unclear whether and how these new documents map to existing documentation that RI organizations already maintain. RI organizations have done a great deal of work

[10] *See, e.g.*, https://arxiv.org/abs/2305.00945.
[11] This terminology is different from that used in Section 3.5.4.2 "PEP Subcomponent 4.2 – Risk Management Plan," and appears to create a heightened (and possibly duplicative) requirement for cybersecurity risks relative to other project risks.

documenting and formalizing their cybersecurity programs, including many documents that have become standard across the community. For instance, many Major Facilities and Mid-Scale Research Infrastructure projects maintain Master Information Security Policy & Procedures (MISPP) documents, Cybersecurity Program Strategic Plans (CPSPs), along with a number of other policies and procedures (*e.g.*, Incident Response Policies, Acceptable Use Policies). Each of these cover at least some of the same ground as the IAMP. But absent more clear guidance on the overlap between these documents, RI organizations may feel compelled to scrap existing policies, procedures, and plans to to meet the new requirements of this section.

Recommendations:

a) NSF should clarify that existing documentation can be used to satisfy the requirements in Section 5.3.
b) For any new documentation requirements, NSF should clearly state what that documentation requires, and clearly identify whether and how new documentation maps to existing documents commonly used in the community.

Comment #8: NSF should clarify the specifics of the Program Assessment in subsection 5.3.11.

NSF closes draft Section 5.3 with a subsection called "Program Assessment." While we strongly agree that program assessments are essential for an effective cybersecurity program, we think this section would benefit from greater clarity on a few topics. For context, many RI organizations currently engage in a number of activities to evaluate their cybersecurity programs, including internal evaluations, third party engagements with Trusted CI and other cybersecurity assessors, and NSF reviews. Clarifying how these existing activities map to this new section would greatly aid our implementation.

In particular, we would appreciate greater clarity on the following:

1) It is unclear how much flexibility RI organizations have to choose the type of program assessment that makes the most sense for their project.
2) It is unclear whether these program assessments were focused on cybersecurity controls, or if they were intended to include programmatic considerations as well.
3) It is unclear whether NSF program reviews that evaluated the project's cybersecurity were themselves considered program assessments.
4) It is unclear whether these program assessments were limited to projects in the Operations Stage, or if they apply to projects still in the Development, Design, and Construction Stages.

Additionally, the following elements are not clear:

1) The third bullet under the second paragraph asks: "Have any significant, reportable, incidents been avoided?" Practically speaking, we are unsure how to report on incidents that have been avoided. Is this bullet meant to ask if any significant, reportable incidents have occurred?
2) Table 5.3.11-1 (Simplified Sample Posture Rubric) is not discussed in the text of the section. It is unclear if this is intended to represent the format RI organizations should use to communicate their controls implementation, is simply an example of a possible format for

doing so, or something else.

Recommendations:

a) NSF should clarify specifics of the program assessment in subsection 5.3.11. In particular, NSF should clarify how this assessment relates to other assessments RI organizations are currently engaging in.
b) NSF should clearly state that RI organizations have the flexibility to address program assessments by the means and mechanisms that match their resources, needs, and existing good practices.