

Roadmap for Securing Operational Technology in NSF Scientific Research

November 16, 2022

Status: Draft Report v1.0

Distribution: Public

Andrew Adams, Emily K. Adams, Daniel Gunter, Ryan Kiser, Mark Krenz,
Sean Peisert, and John Zage

About the 2022 Trusted CI Annual Challenge Team

The 2022 Annual Challenge team consists of Trusted CI members from Indiana University, Lawrence Berkeley National Laboratory, the National Center for Supercomputing Applications, and the Pittsburgh Supercomputing Center.

About Trusted CI

The mission of Trusted CI is to lead in the development of an NSF Cybersecurity Ecosystem with the workforce, knowledge, processes, and cyberinfrastructure that enables trustworthy science and NSF's vision of a nation that is a global leader in research and innovation.

Acknowledgments

In support of this effort, Trusted CI gratefully acknowledges the input from the individuals connected with the following NSF Major Facilities and NSF personnel who contributed to this effort:

- IceCube Neutrino Observatory:¹ Ralf Auer, Steve Barnet, Francis Halzen, Kael Hanson, Benedikt Riedel
- NOIRLab:² Mike Fleming, Chris Morrison, Rod Rutland
- Ocean Observatories Initiative:³ Jeffrey Glatstein, Paul Matthias, Craig Risien, Christopher Wingard
- United States Academic Research Fleet:⁴ Pam Clark, Rose Dufour, Lee Ellett, Ken Feldman, Erich Gruebel, John Haverlack, Jim Holik, Robert Kamphaus, Jon Meyer, Brian Midson, Chris Romsos, Rob Sparrock (ONR), Laura Stolp, Daryl Swensen, Kevin Walsh
- United States Antarctic Program:⁵ Timothy McGovern, Jonathan Prince

We are also grateful to Jim Basney and Kelli Shute of Trusted CI who offered feedback on earlier versions of this report.

This document is a product of Trusted CI. Trusted CI is supported by the National Science Foundation under Grant #2241313. For more information about Trusted CI, the NSF Cybersecurity Center of Excellence, please visit: <https://trustedci.org/>. Any opinions,

¹ IceCube Neutrino Observatory <https://icecube.wisc.edu>

² NOIRLab <https://www.noirlab.edu>

³ Ocean Observatories Initiative: <https://oceanobservatories.org>

⁴ University-National Oceanographic Laboratory System: <https://www.unols.org>

⁵ United States Antarctic Program: <https://www.usap.gov>

findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Using & Citing this Work

This work is made available under the terms of the Creative Commons Attribution 3.0 Unported License. Please visit the following URL for details:

https://creativecommons.org/licenses/by/3.0/deed.en_US

Cite this work using the following information:

Andrew Adams, Emily K. Adams, Daniel Gunter, Ryan Kiser, Mark Krenz, Sean Peisert, and John Zage. “Roadmap for Securing Operational Technology in NSF Scientific Research,” November 16, 2022. DOI: 10.5281/zenodo.7327987 <https://doi.org/10.5281/zenodo.7327987>

Executive Summary	4
1. Summary of Recommendations	5
2. Background and Introduction	6
3. Why is Operational Technology Different?	7
4. Findings Summary	11
5. Roadmap	13
5.1 Mission	14
5.2 Organization & Governance	16
5.3 Policies	20
5.4 Device Procurement & Maintenance	21
5.5 Technical Safeguards and Controls for OT Infrastructure	27
6. Resources	33
7. Next Steps	35
Appendices	37
Appendix 1: The Trusted CI Framework’s Foundation	37
Appendix 2: Summary of Short-Term and Long-Term Recommendations	38

Executive Summary

In 2022, Trusted CI surveyed the practices of National Science Foundation (NSF) Major Facilities with respect to securing operational technology. *Operational technology (OT)* encompasses broad categories of computing and communication systems that in some way interact with the physical world. This includes devices that either have *sensing* elements or *control* elements, or some combination of the two. We consider the term *operational technology* to be interchangeable with *cyber-physical systems (CPS)*. The two tend to be used in the same way but by different communities. Both OT and CPS also encompass *industrial control systems (ICS)*, *supervisory control and digital acquisition (SCADA)*, *Internet of Things (IoT)*, and *Industrial Internet of Things (IIoT)*. OT typically has the capability to be networked but may or may not be actually connected to a network at all times or at all.

Most NSF Major Facilities exist to enable the generation of new knowledge through the operation of scientific instruments at a large scale. These instruments, and the data they produce, are a core component of the NSF Major Facilities’ ability to achieve their missions. The OT that enables these instruments to function is critical to the missions of these Facilities.

This document describes a roadmap that NSF Major Facilities and NSF might draw upon to improve the cybersecurity of their operational technology. It also describes steps that NSF can take to provide more comprehensive and consistent guidance on OT cybersecurity in the Research Infrastructure Guide (RIG)⁶ and other documentation used by NSF Major Facilities for their design and operation. Our roadmap contains both short-term and long-term recommendations and actions. The short-term actions are ones that have the potential to be implemented quickly — within the next 1-2 years. Longer-term actions might take years of planning and not be possible to fully implement until a life cycle refresh of the facility.

Note that while the intended organization to act on those recommendations is, in most cases, an NSF Major Facility (or perhaps NSF Major Facilities and Trusted CI working in concert), the audience also includes NSF itself. This is primarily for awareness but in one case we explicitly recommend action directly by the NSF.

In this roadmap we leveraged the Trusted CI Framework and referenced the NSF Research Infrastructure Guide. In particular, we highlight the “**Musts**” from the Framework which are applicable to the recommendations that we are making to highlight the importance of that recommendation.

1. Summary of Recommendations

The operation of Operational Technology (OT) is central to the function of most major facilities. Cybersecurity is an *enabling* capability to ensure the safe and proper operation of that OT. As a result, in the short term, Trusted CI recommends that the *implicit* centrality of OT security to the scientific research missions of most NSF Major Facilities be made *explicit* in the missions of those Facilities via *explicit* mention in all major organizational statements and via resource allocation to cybersecurity of OT. In the longer term, we recommend that cybersecurity be made central in the charter of NSF Major Facilities in a way that is emphasized by NSF.

⁶ NSF Large Facilities Office, *Research Infrastructure Guide (RIG)*, NSF 21-107, December 2021. https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf21107&org=CMMI.

Trusted CI also recommends that NSF Major Facilities leverage their host institution's security via specialized OT security support where possible, but in the longer term should seek to create an IT/OT specialist role to deal with the extra complexities involved in IT/OT cybersecurity. This role would help bridge the cultural and organizational divide between the portions of an organization currently responsible for IT and OT. Where possible, standardization of equipment (e.g., instruments, research vessels) and configurations both within and between Facilities may also help simplify cybersecurity management within each facility.

Each facility should have well-documented policies in place to identify and mitigate cybersecurity risks to OT. Policies must not only be written but also enforced and maintained over time.

OT security policies should also address security requirements used during both procurement and acceptance testing of OT assets. Procurement and testing should therefore involve both relevant IT security and OT operations personnel. Once acquired, each facility should keep a current inventory of OT equipment.

While many OT security issues are socio-technical or purely social, technical controls must also be in place to enable secure functionality without impeding necessary access within a facility and even outside that facility (e.g., for vendors or remote scientists). This requires understanding the degree of trust that each device must have in other devices on a network to which it is connected. Where an acceptably high level of trust between a set of devices is not present, segmentation and isolation are useful mitigations. We note that segmentation must be consistent to be effective: a firewall that makes exceptions for vulnerable services is not providing much protection.⁷ In the long term, most NSF Major Facilities would be well served to move toward “zero trust” architectures, which seek to eliminate any implicit trust between devices by authenticating each device on the network and isolating it from any other device.

2. Background and Introduction

In 2022, Trusted CI conducted a focused study on the security of *operational technology* used in National Science Foundation (NSF)-funded scientific research. The goal of this year-long effort, involving seven Trusted CI members, was to understand the state of the security of operational technology in science and then to develop a roadmap — this

⁷ Sean Peisert, Matt Bishop, and Keith Marzullo, "What Do Firewalls Protect? An Empirical Study of Firewalls, Vulnerabilities, and Attacks," UC Davis CS Technical Report CSE-2010-8, March 2010. <http://www.escholarship.org/uc/item/9r06p21c>.

document — of clear, actionable recommendations toward sustainable improvement of the security of that operational technology. This roadmap document follows our *Findings of the 2022 Trusted CI Study on the Security of Operational Technology in NSF Scientific Research* report, published earlier in 2022.⁸

The intended audience for this document includes Trusted CI itself, so it can best support the NSF Major Facilities in securing their scientific Operational Technology (OT), those operators of cyber-physical systems in science, and also NSF Program Officers, so that they understand those gaps in securing OT in science and can better understand the need for prioritization and commitment of resources to improving the state of securing OT in science.

Those NSF Major Facilities that have higher degrees of interactions with U.S. Government agencies outside of the National Science Foundation may be required to follow rules pertaining to the regulatory authority of those other agencies (e.g., the U.S. Coast Guard for vessels and marine facilities; the Federal Aviation Administration for aircraft and aviation facilities; and the U.S. Department of Health and Human Services for anything falling under HIPAA regulations). We believe that this roadmap can even help Facilities interacting with agencies with such regulatory authority, although those Facilities will also likely have additional guidance that will be needed to support the additional regulations.

This document was written by team members of Trusted CI, the NSF Cybersecurity Center of Excellence. The team includes security experts from various parts of the discipline including operational security, scientific infrastructure development, and security research.

3. Why is Operational Technology Different?

Operational technology predates computer networks by a very long time and has a long history of being operated safely. In fact, there is an entire discipline called *safety engineering* that seeks to assure that proper engineering principles are used to ensure safety of individuals, equipment, and materials surrounding the use of OT, leveraging fault tree analysis, and ensuring proper failure modes (e.g., fail safe, fail fast, fail slow).⁹

⁸ Adams, Emily K., Gunter, Daniel, Kiser, Ryan, Krenz, Mark, Peisert, Sean, Sons, Susan, & Zage, John. (2022). “*Findings of the 2022 Trusted CI Study on the Security of Operational Technology in NSF Scientific Research (1.0)*”. Zenodo. <https://doi.org/10.5281/zenodo.6828675>.

⁹ C. McParland, S. Peisert and A. Scaglione, "Monitoring Security of Networked Control Systems: It's the Physics," in *IEEE Security & Privacy*, vol. 12, no. 6, pp. 32-39, Nov.-Dec. 2014, doi: 10.1109/MSP.2014.122.

Very large amounts of OT in use today predate the use of modern computer networking and the Internet. Not only was OT *not* designed with network security in mind, it was not designed with networking in mind at all. Networking came later, and with it, attack surfaces of exposed OT on those networks. While it is also true that certain roots of today's computer operating systems also predate the Internet, including UNIX, Microsoft Windows, and Mac OS, all of those systems have seen robust improvements in security over the past decades, whereas the same is not true for a great deal of OT and the software and firmware that controls it.

While operating systems and software written for traditional *information technology (IT)* purposes tend to be general purpose, software and firmware for OT is often custom-written according to the requirements of the device and the function it performs. For the manufacturer of the device or machinery, the focus is usually more on the physical functionality than its software functionality. In fact, companies developing such devices may have created an OT device before the need for software control came about. For instance, a company that makes a winch or crane for a ship may be using basic electronic interfaces without the need for more sophisticated integrated computing or networking. Later, due to demand from customers, the company may implement functional control of the winch or crane via a remote wireless controller or smartphone. While the company may have hired a systems expert to implement the functionality, that expert may not be provisioned to support the software infrastructure over time. Thus, vulnerabilities discovered in the software or firmware may unwittingly go unpatched by the company who may not even be aware that the system is still in operation or should be patched.

By its definition, OT interfaces with the physical world — either via control systems or sensors or both — and may be capable of affecting the physical world. This could include activities that put operators and other personnel in danger of being pinched, crushed, hit, electrocuted, and so on. In the scientific research world, an example of this is in a control system that needs to move a large telescope physically. If the telescope were to move unexpectedly, it could crush an obstruction in its path, including a person. This is in contrast to software that only deals with data and aspects of the virtual world, where physical safety is typically not a direct concern.

OT often can be easy to overlook, as its focus on operations often leads to it working in the background and blending in with the environment. A well-functioning system can easily be ignored. For instance, one may not consider that an heating ventilation air conditioning (HVAC) system in a commercial building (or on a ship) has a network connection for monitoring and control, but that is exactly how the company Target was breached in 2013.¹⁰ Likewise, in a data center, an uninterruptible power supply (UPS) system for power backup

¹⁰ Brian Krebs, "Target Hackers Broke in Via HVAC Company," February 5, 2014. <https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/>.

may have a monitoring and control interface connected to a management network but is rarely utilized until there is an emergency. As with HVAC, vulnerabilities in networked UPS systems were exploited in March 2022 at APC.¹¹

The reality is that there are many vectors into OT systems. Infamously, the Stuxnet malware that manipulated programmable logic controllers (PLCs) to cause uranium enrichment centrifuges in Iran to tear themselves apart jumped network air gaps by being introduced to the facility through a compromised USB drive.¹² A single compromised password enabled attackers to shut down the Colonial Pipeline fuel distribution system.¹³ In 2017, as a result of the NotPetya malware, (1) the radiation monitoring system at the Chernobyl nuclear power plant in Ukraine was disabled; (2) production was halted at Cadbury's chocolate factory in Hobart, Australia; and (3) the shipping giant Maersk's entire computer network, including port operations, was shut down.

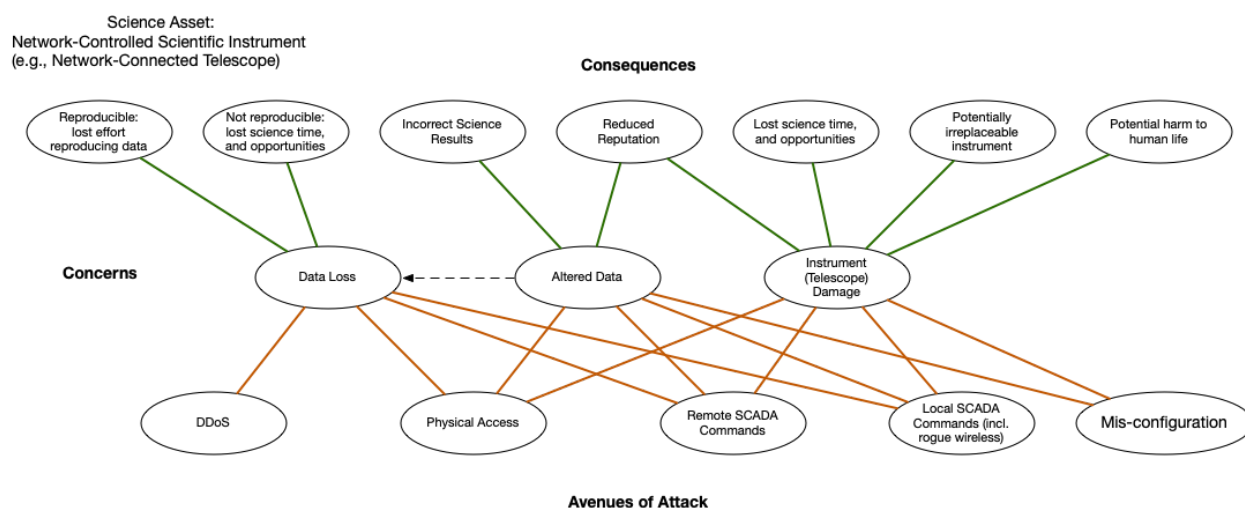


Figure 1: *Open Science Cyber Risk Profile (OSCRP)*¹⁴-style diagram of high-level avenues of attack against network controlled scientific instruments.

¹¹ Eduard Kovacs, "Millions of APC Smart UPS Devices Can Be Remotely Hacked, Damaged," *SecurityWeek*, March 08, 2022.

<https://www.securityweek.com/millions-apc-smart-ups-devices-can-be-remotely-hacked-damaged>.

¹² Kim Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon," *Wired*, Nov. 3, 2014. <https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/>.

¹³ Kim Lyons, "Hackers reportedly used a compromised password in Colonial Pipeline cyberattack," *The Verge*, June 5, 2021.

<https://www.theverge.com/2021/6/5/22520297/compromised-password-reportedly-allowed-hackers-colonial-pipeline-cyberattack>.

¹⁴ *Open Science Cyber Risk Profile (OSCRP)*. <https://trustedci.github.io/OSCRP/>.

The culture and traditions of many organizations often delegate the responsibility of installing, configuring, and maintaining OT to non-IT personnel, vendors, or third-party contractors. These may include contractors brought in to do one job without an obligation to provide ongoing support. There is no guarantee that third-parties have adequate cybersecurity training or awareness of threats, the expectation or motivation to properly secure a networked device, or adhere to fundamental cybersecurity practices such as strong passwords and per-user access controls. Contractors and other third parties may not be aware of the project's security policies. As a result, those individuals may be exposing OT control systems to an organization's IT network or even the broader Internet with minimal security protections in place. There may even be a misalignment of incentives to secure OT systems between contractors and an organization's OT operators and cybersecurity staff. All this results in an increased risk landscape susceptible to successful cyberattacks.

OT deployments are often on separate networks and in unique physical locations, separated from traditional IT devices that might be found in a server room or network closet. Despite this segmentation, and due to their operational requirements, OT might be more exposed to public access, such as a security badge reader or heating and cooling systems. OT may be deployed in physically difficult to access locations to facilitate specialized research, such as at the peak of a mountain or in the depths of the ocean. This isolation or purpose might also require using less secure networks or exposing them on the public Internet to facilitate remote access. Some OT devices require third-party personnel outside of the organization have administrative or operator access to the device to perform maintenance, which in turn may expose the device to external sources, thereby increasing the risk of exposure to network-based malicious activity.

OT devices are often running unique firmware or software that was developed for the specific purpose of the device. This means that if the vendor producing the software is out of business, there may no longer be support for the software even during the expected life cycle of the asset. This is in contrast to more traditional computers where the operating system (OS) was written by one of the major providers with a plan for providing security patching support and a long-term upgrade process.

OT deployment life cycle maintenance and upgrade regiments differ greatly from traditional IT. An additional factor complicating the maintenance of OT device operation is often associated with more expensive equipment than mainstream commodity computers, which are only meant to last five years before being upgraded. The organization's budget might only expect to buy the OT equipment once during the lifetime of the project. Therefore, OT often has different mechanisms for upgrading software or firmware of the device than traditional IT. In fact, the primary function of the device which a facility relies upon may need to be taken out of service for a significant amount of time in order to perform maintenance and updates. Furthermore, when OT equipment is able to receive

security updates, the manufacturer will often only provide security updates on newer equipment. As a result, projects will often run equipment for decades without being able to patch security vulnerabilities. When updates are available, close coordination with the manufacturer or third-party vendor may be required, or possibly an additional purchase that must pass through an organization's procurement process.

4. Findings Summary

For a complete description of our findings on the state of the security operational technology in science, please see our *Findings of the 2022 Trusted CI Study on the Security of Operational Technology in NSF Scientific Research*. Here, for context, we summarize the major findings.

The predominant finding across all NSF Major Facilities we interviewed was that *safety engineering* is taken extremely seriously. Safety engineering is the aspect of operational technology that largely addresses safety of devices that are not typically connected to other networks or other computerized devices (*e.g.*, environmental control systems, fire suppression systems, or winches). Whenever potential damage to equipment or environment or safety of life was at question, OT operators at Facilities had a rigorous understanding of risks involved and policies and procedures in place to address them. However, as we will discuss later in this document, there were consequential gaps when cybersecurity and safety engineering overlapped.

Cybersecurity contains many highly technical elements, and our findings raised concerns about the fact that many operators of OT rely on technical controls such as segmentation or simply disconnection in order to isolate and protect OT. However, we observed instances in which devices periodically need to be connected to the network for updates, thereby nullifying the protection provided by isolation. Moreover, there tended to be limited monitoring of systems — either within a network (“east/west”) or between networks (“north/south”) even when such reconnections occurred.

However, our most significant findings were not technical but unearthed gaps in organizational elements that have potentially an even greater bearing on securing OT. For example, we discovered that despite the outsized risks posed to the missions of NSF Major Facilities by cyber attacks against OT, the portions of the Facilities that operate OT are often disconnected from IT security. While all the Facilities that we spoke with took cybersecurity seriously, and although each facility had one or more individuals responsible for cybersecurity operations, that person was typically from within the IT portion of an organization and was typically siloed in a different part of the organization from OT operations. It is this divided responsibility where security gaps are most prominent, as OT

assets become a source of risks which are unaccounted for by the very parts of the organization whose roles are to account for and remediate cybersecurity risks to the organization's ability to carry out its mission. To this end, every facility indicated that if there was one element of their organization that they could change, it was that they would have at least one FTE specialist dedicated to cybersecurity (including OT security) independent of other responsibilities.

We also discovered that governance of OT cybersecurity is not dictated by NSF and the treatment of which can vary highly between each Facility and even within units or projects at the Facilities (*e.g.*, different capabilities and purpose of telescopes in the NOIRlab Facilities, a variety of form and function of vessels across the U.S. Academic Research Fleet). Conversely, variation was found to be an asset to a facility in that it enables customization of security to the needs of the organization. At the same time, certain degrees of standardization could help organizations with limited resources — as all of the NSF Major Facilities are — by not forcing each organization to reinvent approaches from scratch.

Variation within individual Facilities was particularly prominent when research conducted required physically distributed installation or remote collaborations. A common theme across all Facilities was a lack of available cybersecurity expertise on site at remote locations. For example, there may simply not be enough space (*e.g.*, bunks on a ship) at a remote location for a cybersecurity specialist to be on site at all times. One NSF Major Facility we interviewed mentioned that they have outsourced their cybersecurity expertise as a way to address skill and personnel gaps. This includes outsourcing the security role itself and by leveraging community resources like the Research Security Operations Center (ResearchSOC).¹⁵

Since the NSF Research Infrastructure Guide does not prescribe top-down governance of cybersecurity for NSF Major Facilities, this means that each Facility is left to its own to determine operational policies, including cybersecurity. This includes not just the creation of policies, but the documentation and enforcement of those policies. However, in general, we found very low amounts of documentation and use of OT-related security policies, with a few exceptions relating to NSF Major Facilities that have higher degrees of interactions with U.S. Government agencies outside of the National Science Foundation, or in some way related to export controlled areas. For example, in our interviews, we found policies applied across the Academic Research Fleet were based on policies laid out in the Woods Hole

¹⁵ Research Security Operations Center (ResearchSOC). <https://researchsoc.iu.edu>

Oceanographic Institution (WHOI)¹⁶ Safety Management Manual (SMM)¹⁷ addresses shipboard cybersecurity regulatory requirements.

Contrary to the rapid lifecycle of IT devices and support, facility devices are expected to sustain a number of years or decades of service. Some Facilities are taking a proactive approach in planning for the limited windows of opportunity during facility construction or refresh to establish or enhance cybersecurity in their OT infrastructure and deployments. However, the security properties of either commercially produced OT or “bespoke” scientific OT devices tends not to be well understood by NSF Major Facilities, and nor is security an element of OT procurement requirements, as it might be if it were a traditional computing or networking product or service. Indeed, more than one Facility expressed concern with vendor transparency regarding cybersecurity practices, such as those relating to firmware updates or remote access.

5. Roadmap

In conjunction with our findings report, the Trusted CI OT security study team has divided its recommended roadmap for a path forward for more secure scientific OT into five categories: (1) mission, (2) organization and governance, (3) policies, (4) device procurement and maintenance, and (5) command and control security. The following contains subsections describing each of these categories. In addition, we provide both short and long term recommendations for future mitigation. These recommendations are summarized in **Appendix 2**. Note that the intended organization to act on those recommendations is, in most cases, an NSF Major Facility, or perhaps NSF Major Facilities and Trusted CI working in concert. However, the audience also includes NSF itself, primarily for awareness, but also in one case, recommendation M.3, we explicitly recommend action directly by NSF.

¹⁶ Woods Hole Oceanographic Institution (WHOI) is a highly prominent oceanographic institution that has numerous ships that are part of the Academic Research Fleet.

¹⁷ Woods Hole Oceanographic Institution (WHOI) Safety Management Manual (SMM) chapter index: <https://www-standby.who.edu/what-we-do/explore/cruise-planning/cruise-planning-before-the-cruise/cruise-planning-policies-required-reading/safety-management-manual>.

5.1 Mission

Motivation & Rationale

The *Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators* (Must 1) states:

*“Cybersecurity is not undertaken as an end unto itself: the ultimate goal of a cybersecurity program is to support the organization’s mission. ‘The mission’ is the foundational motivating force driving decision making: it is made up of the task(s), purpose(s), and related action(s) that the organization treats as most important or essential. The program’s implementation must account for the positive and negative impacts security can have on the organization’s mission.”*¹⁸

Most NSF Major Facilities exist to enable the generation of new knowledge through the operation of scientific instruments at a large scale. These instruments, and the data they produce, are a core component of the NSF Major Facilities’ ability to achieve their missions. The OT that enables these instruments to function is critical to the missions of these Facilities.

In many cases, failure or mis-operation of a single large asset at a Facility could result in the entire facility, and the science that it supports, being stalled for years due to downtime and damage. Many Facilities represent sole U.S. capacity for certain scientific disciplines and depend largely on the availability of their OT assets in order to perform their core activities. Consider, for example, the collapse of the Arecibo Telescope in late 2020. The facility has been reduced to a skeleton staff and the science performed by Arecibo is not taking place. Even temporary disruptions in availability could jeopardize the scientific mission. Similarly, the loss — even temporary — of the U.S. Academic Research Fleet’s (ARF) sole icebreaking vessel for arctic research, Research Vessel (R/V) *Sikuliaq*, would have significant impact on Arctic research. Even aside from asset damage, consider the potential consequence of a major safety failure that led to loss of life: an automated telescope control that crushes a human or a ship’s underwater sonar engaged while divers are in the water.

¹⁸ Jackson, Craig, Cowles, Bob, Russell, Scott, Adams, Emily K., Kiser, Ryan, Ricks, Ranson, & Shankar, Anurag. (2021). The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators (1.0). Zenodo. <https://doi.org/10.5281/zenodo.4562447>.

Short-term Recommendations/Actions

M.1. Trusted CI recommends that during NSF Major Facilities’ process of addressing the Framework’s [Must 1](#), they must account for OT assets as well as IT assets within their strategic security plans because of the close relationship between the mission and the OT assets employed in supporting the mission.

This includes acknowledgement and acceptance that any notion that “the science” and “OT security” are not zero sum functions, but that the success of the latter is central to the success of the former. ([Must 1](#)) NSF Major Facilities missions, which typically reflect the use of instruments for scientific discovery, will be the engines that also drive strategic planning as well as the allocation of resources to support the operation of those instruments. Therefore, NSF Major Facilities should explicitly mention cybersecurity of OT in all major organizational statements and also via explicit resource allocation to cybersecurity of OT. NSF Major Facilities may reference the Trusted CI Cybersecurity Program Strategic Plan Template for guidance in prioritizing and formalizing the direction of their cybersecurity program, including the elements involving OT.¹⁹ Of note, Trusted CI offers a number of additional templates and tools relating to establishing and maintaining a cybersecurity program.²⁰

M.2. Strategic plans and budgets must *explicitly* include adequate resource allocations to the cybersecurity of OT. This allocation should not be implicitly buried in other IT security or OT operations line items but called out so that it can be clearly considered in risk management discussions among senior leadership and NSF program management. ([Must 12](#)) We note one of the core findings from our study earlier in 2022 is “every facility indicated that if there was one element of their organization that they could change, it was that they would have at least one FTE dedicated to cybersecurity (including OT security) independent of other responsibilities.” ([Must 11](#) and [Must 13](#))

More generally, as we will discuss in the next section (Organization & Governance), a facility’s mission and strategy require people who are in roles that have responsibility for cybersecurity both as a whole and for OT specifically, and likewise have the authority to make decisions with regard to cybersecurity risk acceptance or mitigation. ([Must 6](#))

Long-Term Recommendations/Actions

M.3. NSF should emphasize the centrality of OT cybersecurity in charters of NSF Major Facilities. The current language only mentions OT security in the context of

¹⁹ Trusted CI. Template: Cybersecurity Program Strategic Plan V1.0.

<https://docs.google.com/document/d/1Z6SkgxVFGi-Aw4oSaRHETmoqfpGygeBb5kAvNykQYjQ>.

²⁰ Trusted CI Framework Templates and Tools. <https://www.trustedci.org/framework/templates>.

ensuring technical controls around Industrial Control Systems (ICS) and Supervisory Control and Data Acquisition (SCADA) systems, and not as part of the broader guidance around establishing and maintaining a robust cybersecurity program. We believe that this is insufficient. A starting point for revising the charters of NSF Major Facilities is for NSF to more fully integrate OT language into the Research Infrastructure Guide (RIG), which itself points to the Trusted CI Framework for guidance structuring and implementing a cybersecurity program. NSF itself should emphasize this guidance in cybersecurity for *both* OT and IT when awarding Major Facilities.

M.4. NSF Major Facilities should work both inside their organizations and outside of them, with other Facilities, to try to standardize IT and OT configurations.

In addition, though instrumentation in NSF Major Facilities can often be both bespoke and unique — as we have indicated, Facilities can represent sole source instrumentation capability for a scientific discipline — where possible, standardization of equipment (*e.g.*, instruments, research vessels) can help by developing boilerplate solutions that can be consistently used across Facilities. Standardization of OT can help reduce the need for dedicated OT FTEs per site, because they can then be shared across sites.

M.5. NSF Major Facilities should work with Trusted CI and CI Compass on coordination across facilities. Trusted CI and CI Compass²¹ are both NSF-supported Centers of Excellence whose missions are to support cyberinfrastructure of NSF Major Facilities. Both are already actively involved in helping NSF Major Facilities coordinate between each other and are ideally positioned to support the coordination of the secure operation of OT among Major Facilities as well.

For more discussion of Mission Alignment, please refer to the Trusted CI Framework Implementation Guide “Mission Alignment” section, which includes [Must 1](#) and [Must 2](#).

5.2 Organization & Governance

Motivation & Rationale

Organization & Governance (O&G) includes organizational guidelines and review procedures for assuring secure and safe operation. It plays an important role in long-term planning and decisions about resource allocation for both people and material. Findings from the 2022 Trusted CI Study on the Security of Operational Technology in NSF Scientific

²¹ CI Compass, the National Science Foundation Cyberinfrastructure Center of Excellence (CI CoE). <https://ci-compass.org/>.

Research²² give us evidence of which sections of the Trusted CI Framework are needed for improving O&G. The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators covers relevant topics in [Musst 7, 8, 13, and 14](#).

We observed through our interviews with the subset of facilities that they appear to have effective cybersecurity leadership. However, from input from the Framework Cohort, many institutions do not have an assigned cybersecurity leadership role or have appointed one from among the staff available, without regards to cybersecurity expertise. Following the guidance put forth in [Must 7](#) of the Framework Implementation Guide, NSF Major Facilities must establish a clear lead for their cybersecurity responsibility:

“Organizations must establish a lead role with responsibility to advise and provide services to the organization on cybersecurity matters. Due to the complexity and breadth of cybersecurity issues and the need for coordinated decision making, organizations require an individual role to lead cybersecurity.”

For a cybersecurity program to be effective, it needs to reach remote locations and their personnel. This is mentioned in the Framework Implementation Guide in [Must 8](#): Comprehensive application, “Organizations must ensure the cybersecurity program extends to all entities with access to or authority over information assets.” As mentioned in the finding’s report section 3.2, there is frequently “siloeing” between information security personnel and OT operators. This siloeing prevents effective cooperation between the cybersecurity program and OT operators.

For Major Facilities, there is a need for qualified personnel for handling the complexities of OT and security related to it, as mentioned in our Findings report: “host institutions often do not have the specialized skillset or personnel available to support OT security.” Also as mentioned in [Must 13](#): Personnel in the Framework Implementation Guide:

“Personnel resources are commitments made by an organization to assign human effort to particular activities on behalf of the organization. Personnel resources allocated to cybersecurity include both full-time cybersecurity employees and employees with partial cybersecurity responsibilities. Personnel resources allocated to cybersecurity may be assigned to carry out a number of organizational activities, including security operations, governance, management, architecture, and incident response.”

²² Emily K. Adams, Daniel Gunter, Ryan Kiser, Mark Krenz, Sean Peisert, Susan Sons, and John Zage. Findings of the 2022 Trusted CI Study on the Security of Operational Technology in NSF Scientific Research, Trusted CI, July 13, 2022. DOI: 10.5281/zenodo.6828675. <https://zenodo.org/record/6828675#.Y1ftjezMJnk>.

Making the situation more complicated, assessments of the current personnel already hired are needed to ensure their awareness of and ability to perform tasks for cybersecurity.

With the difficulty with having the required experts for both IT, OT, and cybersecurity, [Must 14](#) (External Resources) points to how to cover the gap if finding internal experts doesn't work out:

“External resources include services, tools, and collaborators outside of the organization that can be leveraged to support the cybersecurity program. Identifying them, picking judiciously, and using them can greatly benefit the organization and optimize local resources. Because the external organizations vary widely, leveraging these resources requires careful, advanced planning to maximize the benefit to the organization.”

Also there is a need for external resources for maintaining OT devices, as mentioned in the findings report: “Some OT devices require that third-party personnel outside of the organization have administrative or operator access to the device to perform maintenance, which in turn may expose the device to external sources, thereby increasing the risk of exposure to network-based malicious activity.”

Short-term Recommendations/Actions

OG.1. Appoint a cybersecurity lead with oversight over both IT and OT security, if one is not already in place. Having one individual appointed the key role for cybersecurity responsibility for both IT and OT is necessary for organizing and maintaining a cybersecurity program. ([Must 7](#))

OG.2. MFs should collaborate with their host institution's existing IT and cybersecurity support organizations to help them to address their OT security needs. While their host's security offerings may not be a perfect fit, they can offer stop-gap measures until the MF's security program is more mature. ([Must 14](#))

OG.3. Assess the degree of on-site OT expertise and their awareness of cybersecurity topics, especially at remote locations. ([Must 8](#))

OG.4. Ensure the representation of OT personnel within institutional cybersecurity groups. ([Must 8](#))

OG.5. Develop guidelines for outsourcing expert cybersecurity roles of personnel and find additional guidance to ResearchSOC²³ and/or commercially supported equipment when in-house resources are not available. The Research Security Operations Center (ResearchSOC) at Indiana University helps make scientific computing resilient to cyberattacks and capable of supporting trustworthy, productive research. It does this by providing the operational cybersecurity services, training, and information sharing necessary to a community as unique and variable as research and education (R&E), and can support NSF Major Facilities by adding OT cybersecurity expertise where it is otherwise infeasible for a Facility to hire such an individual directly. (Must 14)

Long-Term Recommendations/Actions

OG.6. Provision a Chief Information Security Officer or cybersecurity lead role if the current role is not already filled with an individual whose sole role is cybersecurity and has the expertise to fulfill it. (Must 7)

OG.7. Provision an IT/OT specialist role to deal with the extra complexities, a role for cybersecurity IT/OT experts, or a combination of the two. The need for both IT/OT experts and cybersecurity experts may be combined depending on the difficulty of having too many personnel at remote locations, but that would increase the difficulty of finding qualified individuals. (Must 13) See, Section 7 (“Resources”) about the NIST National Initiative for Cybersecurity Education (NICE) framework for more information.

OG.8. Creating a long term institutional goal for IT/OT collaboration and communication on top of the short term representation of IT/OT personnel is essential for buy-in to cybersecurity goals of the organization. Organizational culture may need to be adjusted for cybersecurity to be accepted as an aid to performance, rather than a necessary evil that hinders a worker’s effectiveness. These changes would allow for a more comprehensive application of the cybersecurity program. (Must 8)

²³ Research Security Operations Center (ResearchSOC). <https://researchsoc.iu.edu>.

5.3 Policies

Motivation & Rationale

The Trusted CI Framework Implementation Guide provides the following description for [Must 9](#):

“Policy’ refers to documented normative statements adopted by an organization to govern human behavior. These include authoritative documented statements of “policy,” but can also include “procedures” and other normative guidance. Some amount of policy is needed to formalize and communicate about a cybersecurity program. Processes to develop, adopt, explain (e.g., , provide notice and training), follow, enforce, and revise policies are necessary to make policies an effective component of a cybersecurity program, and keep the policies in line with the organization’s mission.”

Policies should include onboarding and offboarding procedures, procedures for determining escalated privilege roles and durations, a point-of-contact incident response database for each escalated privilege role, and provide for periodic audits of current roles/capabilities. Most importantly, though, the policies need to be approved by leadership, ensure that all personnel are familiar with the processes, and that exceptions are handled appropriately — exceptions, although necessary, should be rare.

Short-term Recommendations/Actions

P.1. Include a review of OT infrastructure as part of an organization's regular assessment of assets. Part of a recommended annual review of cybersecurity policies should include a focus on how they apply to OT assets. When an organization first adopts cybersecurity policies, they may have avoided OT asset considerations, thus it is important to reassess the policies with a different perspective. Also, determine if policies are needed to address risks that cannot be addressed through automated controls ([Must 10](#)).

P.2. Provide policy implementation and awareness training for staff, leadership, and affiliates for policies impacting the NSF Major Facility and research projects. It is recommended that the organization schedule regular meetings (i.e., monthly or quarterly) where cybersecurity staff cover new and review existing policies ([Must 9](#)). Because OT may be viewed by some as outside the purview of a cybersecurity team, the organization should set policy to provide cybersecurity training and awareness about the OT assets for leadership as well as any staff involved in its use or maintenance.

P.3. Require, through policy, that supplemental controls be applied to protect the OT device. For OT that lacks sufficient security controls as referenced by the adopted baseline control set ([Must 15](#)), it will be necessary to seek additional and alternate controls as warranted ([Must 16](#)); one such possible alternate control set or ‘best practices’ guidance is the NIST Special Publication 800-82r3 (see [Section 6](#).) For instance, in the case that a network connected device lacks the ability to protect itself using Access Control Lists, placing a firewall in front of the device that limits access will provide an additional layer of protection (see [Section 5.5](#)). Similarly, scientific research is often pushing the boundaries of technology and projects may be manipulating OT in ways it was not designed to be used. e.g., An underwater drone spoofing GPS in order to get the parent ship's navigation to follow the vehicle. These cases should be cataloged and a plan should be formed to develop controls to reduce cyber risk.

P.4. Establish a policy exception procedure to accommodate OT considerations. Since OT may not be governed completely/adequately by a specific policy, the OT operator may seek an exception for the asset. It is important to remember that exceptions, although necessary, should be rare, and it is recommended that exceptions to policies follow the guidance provided under [Must 9](#) of the Trusted CI Framework Implementation Guide.

Long-Term Recommendations/Actions

P.5. NSF Major Facilities should participate in industry forums to voice concerns over security risks and support efforts to fix such insecurities. Some aspects of OT security deal with inherent insecurities in the technology (e.g, GPS). In order to help solve these bigger security issues, it is recommended that where there exist alternative technologies that are more secure and satisfy the organization's needs, OT operators should consider migrating to those as a long-term goal.

5.4 Device Procurement & Maintenance

Motivation & Rationale

The OT asset procurement processes, the nature of the device deployments, and the ongoing serviceability and maintenance of operational technology components all play a significant role in facilitating scientific research and within organizations’ facility operations. Considerations surrounding acquisition of assets and the lifetime maintenance of these assets are a point-of-origin for an organization's function and can have a lasting effect on research and facility operations as devices may be in production for years or even decades.

The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators [Must 14: External Resources](#) addresses leveraging external cybersecurity resources to support an organization's cybersecurity program. While [Must 14](#) focuses on IT-related solutions and tools provided by external vendors, the following tenets presented in [Must 14](#) hold true for OT procurement and vendor relations:

OT Device Procurement:

“RCOs should also give careful consideration to the following areas when considering products and services: characteristics of the product itself; visibility into activity essential for discovering emergent security concerns before they become bigger problems; visibility into product performance (service or functionality expected from the provider); and resources needed to support the use of the product.”

OT Device and Infrastructure Maintenance:

“It will benefit an RCO to focus on product vendors and services that are already following security best practices relative to their area of focus, are responsive to project-specific security concerns, are communicative, and have procedures in place for mitigating emergent security issues. How the RCO manages these relationships has a substantial effect on the risks and costs associated with information security.”

OT device installations and OT infrastructure can include commercial off the shelf hardware (COTS) and software technologies that are often designed and made in any number of different countries. Acquisition processes can have little, if any, aspect that considers device origin when making purchasing decisions. Likewise, a requirement to assess the cybersecurity capabilities of or cybersecurity protocols used by OT devices are often not built into the selection and procurement process, as there are instances where an OT device cannot, by design, employ cybersecurity protections or cybersecurity protocols in its intended operation.

The quality and consistency of vendor-supplied support for OT devices integrated into research operations can significantly impact the viability of ongoing physical operations, as can *vendor lock-in*: the relationship between a customer and a single service provider where the customer becomes dependent on the provider due to restrictions in service agreements or even in the event a device is so specialized it is only offered by a single vendor. Vendor transparency regarding cybersecurity practices, integrity of device patches and updates, remote vendor access to a device, and avoiding vendor lock-in all contribute to the efficacy of a facility's cybersecurity stance.

Explicitly integrating OT cybersecurity considerations within the existing procurement procedures, vendor relations and accountability, and ongoing OT device maintenance

activities provides further opportunity to protect the organization, facilities, and scientific research from harm.

It is important to note the distinction between a *decision-maker* who determines which OT solutions should be purchased to meet the operational needs of the research function, and the *purchasing entity*, such as the procurement office of said OT solutions. Likely both will have to work closely together, and with the OT security lead to discuss security related issues with procurement.

Short-term Recommendations/Actions

PM.1. Ensure the facility has a current inventory of OT environments and devices. In order to move forward with recommendations related to procurement and maintenance of OT devices and installations, the facility should have full awareness of device types, device purpose, device ownership, and current state of software and/or firmware versions they currently own/employ. [Must 3](#) highlights a number of factors that describe the importance and consequences of documenting organization information assets. If possible, include information that indicates the nature of vendor support (*i.e.* no support, shared support, or vendor-only support). The collection of this information in a central location will better situate the facility to make informed decisions in future procurement and expedite support when interfacing with vendors, especially in situations that require urgent attention (*e.g.*, cybersecurity events or emergency maintenance).

PM.2. Educate both the *decision-makers* and *purchasing entities* on cybersecurity risks within OT solutions. [Must 5](#) describes the importance of involving leadership in cybersecurity decisions, which necessarily requires educating individuals in these roles as to the key considerations. This can be an avenue to mindfully evaluate cyber risk and proactively integrate cybersecurity into facility and research operations. An excellent foundational resource is the *Department of Homeland Security Cyber Security Procurement Language for Control Systems* which “summarizes security principles that should be considered when designing and procuring control systems products and services (software, systems, maintenance, and networks), and provides example language to incorporate into procurement specifications.”²⁴ Both *decision-makers* and *purchasing entities* need to be cognizant of defined security requirements, regulations, and/or periods of compliance associated with OT-related acquisitions and deployments (*e.g.* security requirements for devices acquired with federal funds).

²⁴ Department of Homeland Security Cyber Security Procurement Language for Control Systems. https://www.cisa.gov/sites/default/files/documents/Procurement_Language_Rev4_100809_S508C.pdf.

PM.3. Implement technical controls to secure vendor and third-party access to OT devices. Vendors may require access to OT devices for maintenance, updates, and troubleshooting activities. It is not generally necessary to broaden vendor’s ingress to additional IT or OT systems (see [Section 4.5](#)). Likewise, components and systems shared with groups to facilitate research should be restricted to only those activities necessary to conduct research. Vendor access may not be part of a standard control set, so therefore [Must 16](#) requires additional controls be put in place.

Facilities should evaluate existing protections beyond IT-OT network separation for security. By doing so, Facilities protect themselves against unintentional (or in the unfortunate event, intentional) cybersecurity activities that may have destructive repercussions. *See* section 4.5 for key strategies in “locking down” access to OT devices which include firewalls, network access control, two-factor authentication, and can even include isolating a device from the network which requires physical access to the device. Looking beyond technical security measures, procurement guidance and research agreement language may be used to establish a precedent guaranteeing secure practices like vendor remote connections and shared research access to OT devices.

PM.4. Clearly identify what service-level activities may threaten or negate vendor support and maintenance of OT equipment. As stated in the Rationale section, some activities such as patching or modifying of ICS components without notifying and/or involving the vendor can nullify the system warranty. If the purpose or function of an OT device is to change from its current state or deviate from the original operating intent, validate that these modifications will not disrupt or negate ongoing vendor support. Use this recommendation as a motivator to document and build best practices for OT equipment maintenance and engaging with vendors. *See also*, the long term recommendation regarding an OT acceptance testing program.

PM.5. Begin building security requirements of OT equipment into vendor support contracts and shared research agreements. Clearly identifying, specifying, and/or validating with the OT vendors the cybersecurity-related capabilities can make for a stronger cybersecurity foundation OT operations and infrastructure and better situates the facility to perform reliable and uninterrupted research operations. Validating the cybersecurity requirements of OT products against both current and future threats over the projected product lifespan can be used to inform future decisions about buying products and services. Some key considerations to investigate when selecting devices and entering vendor support contracts are as follows:²⁵

²⁵ Some items in this list have been informed by guidance in NIST Special Publications 800-213 “IoT Device Cybersecurity Guidance for the Federal Government: Establishing IoT Device Cybersecurity Requirements” (<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-213.pdf>), and

- Device and device components origin (geographical)
- Cybersecurity protocols used during device operations and data transfer
- Ability to support local or interfaced device authentication
- Ability to configure aspects of the authentication mechanism
- Ability to restrict unauthorized interactions
- Ability to update and patch the device
- Ability to access threat and vulnerability management options

PM.6. Start integrating language related to cybersecurity requirements into purchasing best practices and/or policies. To facilitate a heightened awareness of cybersecurity protections and to integrate this awareness into practice, the facility can integrate language related to cybersecurity requirements into purchasing best practices or policies ([Section 4.3](#)). Documenting said activities can, at minimum, provide a checkpoint to both the *decision-maker* and the *purchasing entity*. Awareness of these and similar cybersecurity considerations within the existing procurement procedures can protect the facility and scientific research conducted therein. Reiterating recommendation PM.2, both *decision-makers* and *purchasing entities* need to also be aware of externally-defined security requirements, regulations, and/or periods of compliance. Even consider including legal counsel, contract administration, and research project management when developing ongoing security requirements for OT matters.

Long-Term Recommendations/Actions

PM.7. Integrate cybersecurity supply chain risk management (C-SCRM) considerations into organizational policies, plans, and practices. C-SCRM offers an opportunity to establish cybersecurity in the organization’s acquisition ecosystem. Organizations should also familiarize themselves with NIST SP 800-161,²⁶ Supply Chain Risk Management Practices for Federal Information Systems and Organizations [SP800-161]. Organizations should begin, or continue, implementing key practices related to C-SCRM security controls and C-SCRM risk management process.²⁷

PM.8. Clearly communicate expectations and requirements to vendors. Following the previous recommendation, *explicitly* extending the cybersecurity expectations and

800-213A “IoT Device Cybersecurity Guidance for the Federal Government: IoT Device Cybersecurity Requirement Catalog” (<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-213A.pdf>).

²⁶ NIST SP 800-161 Rev. 1. Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. May, 2022. <https://csrc.nist.gov/publications/detail/sp/800-161/rev-1/final>.

²⁷ NIST SP 800-82 Rev. 3 (Draft) Guide to Operational Technology (OT) Security, April 2022. <https://csrc.nist.gov/publications/detail/sp/800-82/rev-3/draft>.

requirements to vendors will (1) help them gain better understanding of local expectations/policy, (2) establish an understanding of a guaranteed secure remote connection between the vendor and the facility's control systems. To help conduct this communication the facility should develop a process to vet suppliers and service providers to ascertain their capabilities, trustworthiness, and the adequacy of their internal security practices.

PM.9. Develop plans-of-action for vendor support limitations. Unfortunately, due to specialized research activities and unique OT device functions required in scientific research at the facility, there is a risk of “vendor lock-in” due to limited availability in the market for specialized equipment. This lock-in can be precipitated by a number of factors including proprietary devices serviceable by a single vendor or a research function that continues operation even after the device vendor has gone out of business. In some cases this is unavoidable, so Facilities should develop plans-of-action to address OT system vendor support limitations. Two examples of such plans-of-action are assessing limitations within bounds of existing service agreements and ensuring internally-sourced support activities (*see*, earlier patch-management short-term recommendation) can be conducted in the event vendor support becomes unavailable. Also, bolstering OT expertise within the facility workforce can mitigate the danger of maintaining and operating OT devices that no longer have vendor support options.

PM.10. Develop and implement an OT acceptance testing program. Acceptance testing is a technique used to determine if systems or components have met requirement specifications, evaluate the system's compliance with the business requirements, and verify the criteria for end users are met. An OT acceptance testing program can (1) be leveraged to validate the secure function of the device upon acquisition and during device patches and upgrades, and (2) be used as a vehicle to set researcher and vendor expectations. As acceptance testing can be an involved process sometimes requiring specialized expertise, consider leveraging peers or host institutions with experience or programs geared towards device and device function validation.

Acceptance testing can also involve lengthy time-driven activities like approval, implementation, and deployment of patches/devices, thus, collaborating with researchers to develop their program timelines and operational expectations is critical. Clearly communicating with vendors about the facility acceptance testing program may prompt increased transparency of patch details and impact to device functionality as information related to patch specifications are often not provided by vendors. Ultimately, an OT acceptance testing program will promote greater visibility into interruptions to OT operations and ultimately situate the facility and researchers to conduct dependable and consistent research activities with their OT devices.

5.5 Technical Safeguards and Controls for OT Infrastructure

Motivation & Rationale

OT networks and the assets on them are not designed to be resilient to unexpected network conditions, let alone deliberate misbehavior targeting these assets. Because of this, these assets must be separated from typical IT networks. *Network segmentation* is therefore a key cybersecurity control. Ultimately, underlying these approaches is a set of principles which we believe align with the concept of *zero trust*. For those familiar with zero trust, it is important to understand that we do not necessarily recommend full-blown zero trust architecture implementations in the near term, and indeed for some types of OT devices, they may not be appropriate at any time in the foreseeable future. With that said, the motivation and principles supporting zero trust are important when considering how networks and segmentation influence technical controls used to safeguard OT.

Zero trust principles and *zero trust architecture* have been defined by various entities. In general, the *principles* center around the notion of “eliminating implicit trust and continuously validating every stage of a digital interaction” and ... “never trust, always verify.”²⁸ These principles have architectural implications at the network level that include “all users, whether in or outside the organization’s network, to be authenticated, authorized, and continuously validated for security configuration and posture before being granted or keeping access to applications and data,” and assuming “there is no traditional network edge [whether] networks [are] local, in the cloud, or a combination ... with resources [and] workers in any location.”²⁹

NIST SP 800-207,³⁰ an authoritative reference on zero trust uses similar definitions for zero trust while also emphasizing the planning required for such an implementation:

“Zero trust architecture (ZTA) is an enterprise’s cybersecurity plan that ... encompasses component relationships, workflow planning, and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan.”

²⁸ What is a Zero Trust Architecture, Palo Alto Networks.

<https://www.paloaltonetworks.com/cyberpedia/what-is-a-zero-trust-architecture>

²⁹ Zero Trust Security Explained: Principles of the Zero Trust Model, CrowdStrike, October 17, 2022.

<https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/>.

³⁰ Zero Trust Architecture, NIST Special Publication 800-207, August 2020.

<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>.

At the same time, SP 800-207 also provides detailed technical strategies that may be opaque concepts to OT operators. Due to these two complexities, we endeavor to build off of the more consumable definition(s) of zero trust and zero trust architecture so that we can provide clear and actionable recommendations based on those principles. In short, we believe OT operators should interpret zero trust as:

Assume your local-area networks are compromised. To protect services and their data running within those networks, ensure that (i) only authorized users (users that satisfy the principle of least privilege and after authentication) are allowed to access said service, and (ii) that only authorized collaborating services can communicate with said services.

Short-term Recommendations/Actions

T.1. OT systems should only run necessary tasks. Baseline configurations for these devices should be strictly limited to the intended operation of the system wherever possible. However, legacy systems may have any number of active but unused services that may not be disabled or blocked. These active unused services and communications ports within the ICS component present a cyber security issue.³¹ (Must 15)

T.2. Document and monitor network and network-attached assets. In order to begin to segment networks it is important to develop sound documentation describing the current state and topology. In particular, it is crucial to develop comprehensive and usable network diagrams. These will be instrumental to the ability of responsible individuals to reason about networks, to train and onboard staff, to identify components which can be better isolated, and identify other opportunities to improve the underlying network architecture. This documentation process should result in two types of network diagrams: a physical topology diagram showing physical cabling and connections between different hardware components; and a logical topology diagram showing which assets should communicate with one another as well as the protocols by which they connect. Physical network diagrams can be important because they can assist with troubleshooting, indicate how networks might be expanded, show how physical connections and disconnections can be made, and show where network monitoring can be integrated if necessary. (Must 3)

Similarly, as important it is to document your network assets, it's equally important to monitor what traffic flows across your networks. Every baseline control set (Must 15) will specify some sort of network monitoring control, for it's essential for administrators to understand which devices are communicating and under what protocols the traffic is being

³¹ U.S. Department of Homeland Security. Recommended Practice for Patch Management of Control Systems. December 2008.
https://www.cisa.gov/uscert/sites/default/files/recommended_practices/RP_Patch_Management_S508_C.pdf.

transported in order to determine if the traffic is both authorized and efficient. Be aware though, that in some low bandwidth or high latency networks, the cost of monitoring may add excessive burden to the network. Thus, at the very least, in resource-limited networks monitoring should have its impact measured prior to being enabled completely.

T.3. Identify network-attached control ports. Additional issues can arise from misuse of computers connected to these networks for control purposes. For example, an open USB port on a computer used to control propulsion or navigation systems on a ship introduces the risk that someone could attach a USB storage device which has been compromised or a cellular phone which is connected to other networks, effectively bridging the security boundary. Organizations should carefully control configurations of PCs and other devices with external interfaces connected to OT networks to prevent this. Baseline configurations for these devices should be strictly limited to the intended operation of the system wherever possible. At minimum, we recommend the following measures (see [Appendix 3](#) for additional details on these recommendations):

1. Document any external interfaces such as USB ports and network interfaces and their expected condition. Review the state of these devices periodically to ensure that the devices do not deviate from this established baseline.
2. Configure systems to ensure records of important events are kept in logs and backup logs to central logging systems.
3. Restrict use of USB storage devices to a set of approved users who need this ability to perform their job responsibilities.
4. Restrict all ability to change any network configuration, including networking configuration, firewall configuration, cellular tethering, and addition of new interfaces on systems to only those staff responsible for configuration and maintenance of these systems.
5. Users of these systems should not be granted elevated privileges without additional authentication steps.
6. Wherever possible, administrative accounts — accounts with elevated privileges — should not be used to perform the day-to-day functions of the systems.
7. Staff with elevated privileges on these systems should be trained to understand the potential safety and security implications of the use of their elevated privileges and the importance of maintaining any established security boundaries such as air gaps.

For human-machine interface (HMI) systems or other embedded systems or controllers with external data interfaces such as USB or serial ports where these additional controls are impossible to implement, physical controls should be put into place to prevent the use of these interfaces by unauthorized individuals. These may include ensuring that these interfaces are disabled, the external interfaces are only accessible behind locked access

ports or doors, or the system itself is in a location which is inaccessible to unauthorized staff members.

T.4 The principle of least privilege should be applied to control access to OT assets. *Identity and access management (IAM)* refers to a process and lifecycle for mapping human beings to computing accounts, determining which resources those accounts should have access to, determining what rights those accounts should have with respect to reach of those resources, and managing those access rights over time as roles, responsibilities, personnel, and resources change. Application of the *principle of least privilege*³² can be performed with IAM solutions *along with* policies for determining what roles and capabilities should be assigned to staff, third-party collaborators, and guests. We note that Trusted CI maintains a large set of resources regarding IAM, including hands-on training, documentation, webinars, a monthly working group, and email discussion lists.³³

T.5. Ensure the ability to recover from disaster. The ability to recover critical systems should be prioritized as well. For OT systems this can be difficult to accomplish without support and intervention from the vendor. For this reason, emergency recovery mechanisms and provisions must be a consideration when making procurement decisions (see [Section 5.4](#)).³⁴ If the vendor does not provide sufficient recovery capabilities, it may be necessary to develop them internally. This may necessitate keeping copies of configurations and firmware as well as necessary connectivity mechanisms and software tools to restore them. If this is necessary, a form of integrity checking should be recorded separately when backups are taken so that it's possible to ensure that the stored configuration is identical to the known-working state at time of backup. ([Must 15](#))

T.6. Collaborating OT services should be segmented. Practically speaking, this can be achieved in many cases by carefully assigning those services to private *virtual local area networks (VLANs)*, where only allowed hosts and ports are allowed to communicate. The most complicated part of this is determining what services need to be allowed to communicate and to what systems. For example, internal services on private VLANs must be able to speak to my DNS service, but the DNS service should *not* be receiving packets from outside (but for zone transfers if it's a secondary). Likewise, OT should only speak to the third party responsible for applying patches/updates, and the hosts and ports that the third party is connecting from need to be known. For example, a printer service should be accessible to any other services in the organization but should not be accessible to anything

³² https://en.wikipedia.org/wiki/Principle_of_least_privilege.

³³ Trusted CI. Identity and Access Management. <https://www.trustedci.org/iam>.

³⁴ U.S. Citizenship and Immigration Services. Handbook for Employers M-274 Section 4.4. <https://www.uscis.gov/i-9-central/form-i-9-resources/handbook-for-employers-m-274/40-completing-section-2-of-form-i-9/44-automatic-extensions-of-employment-authorization-andor-employment-authorization-documents-eads-in>.

outside; moreover, the printer services should not be able to initiate a connection to any services inside or out.

Note, however, that there can be challenges in using VLANs to segment OT. For example, while most modern computing and networking systems assume that an Ethernet frame has a predictable structure that includes provisions for specific elements such as VLAN tags, there exist devices that produce messages that *look* like Ethernet frames, but work on different data structures. In such a case, those systems can break in unexpected or intermittent ways if one were to try to add VLAN tags to traffic sent to those devices.

As a result, the perfect solution for segmenting may not currently exist, or may vary highly, and require careful consideration. For example, perhaps rather than *enforcing* segmentation, it might be preferable to *monitor* network traffic and alert on violations. In these complicated situations, rather than recommend a specific solution in this roadmap, we strongly encourage consideration of the *principles* of zero trust even if a full implementation cannot be achieved. Trusted CI can help NSF Major Facilities weigh the various options.

T.7. Organizations should carefully control configurations of PCs and other devices with external interfaces connected to OT networks. Systems that communicate with OT assets may end up bridging two or more networks that include *both* IT and OT assets. This is undesirable because this system then provides a vector between networks that are otherwise thought to be properly segmented. However, such a scenario might also be unavoidable in some situations. As a result, these systems must be particularly tightly locked down and carefully controlled in terms of access to such systems and access that such systems are permitted. Both must be minimized, as must the number of different services that each system provides.

Long-Term Recommendations/Actions

T.8. Move to a Zero Trust Architecture. In order to protect OT assets effectively, Facilities must adopt a comprehensive strategy for ensuring the segmentation of these networks, the control of data flows between nodes on these networks, the security of IT assets connected to OT networks, and the behavior of the users of these systems.

For existing infrastructure, the adoption of a zero-trust model may be impractical in the near term. These solutions can require large staffing and financial resource commitments to set up and maintain, as well as a sound understanding of the organization's network architecture. Vendor solutions can rely on network capabilities which OT assets sometimes do not support or implement effectively. Despite these issues, a large degree of network

segmentation is very valuable when securing these critical networks. We believe that this model of network segmentation is ultimately desirable for OT networks and resources should be allocated to attempt to move in this direction.

Going forward, NSF Major Facilities should look for vendors producing OT that *does* support the segmentation and authentication approaches to enable a zero trust architecture. And, for bespoke scientific OT, should emphasize to the developers of such assets the need for support of approaches that enable zero trust as well.

6. Resources

Beyond the recommendations made in this document there are additional resources available to help orient NSF Major Facilities to the best practices in OT deployments and organizational considerations for OT. Furthermore, [Must 14](#) of the Trusted CI Framework recommends the use of external resources as part of your overall cybersecurity program.

Trusted CI, the NSF Cybersecurity Center of Excellence

The mission of Trusted CI is to lead in the development of an NSF Cybersecurity Ecosystem with the workforce, knowledge, processes, and cyberinfrastructure that enables trustworthy science and NSF’s vision of a nation that is a global leader in research and innovation. From quick questions to collaborative engagements lasting months, Trusted CI tackles challenges of all sizes. Contact Trusted CI at info@trustedci.org

National Institute of Standards and Technology (NIST)

NIST has produced a Special Publication 800-82r3 draft document “Guide to Operational Technology (OT) Security” which “provides guidance on how to secure operational technology (OT), while addressing their unique performance, reliability, and safety requirements.” The document provides an overview of OT and typical system topologies, identifies typical threats and vulnerabilities to these systems, and provides recommended security countermeasures to mitigate the associated risks.³⁵ NIST has also developed specific guidance for the application of the security controls to OT in NIST Special Publication (SP) 800-53 Revision 5 “Security and Privacy Controls for Information Systems and Organizations”, which is included in Appendix F of the document.³⁶

Guidance contained in NIST Special Publications 800-213 “*IoT Device Cybersecurity Guidance for the Federal Government: Establishing IoT Device Cybersecurity Requirements*”³⁷, and 800-213A “*IoT Device Cybersecurity Guidance for the Federal Government: IoT Device Cybersecurity Requirement Catalog*”³⁸ offer key considerations in the management of operational devices and deployments. While the titles indicate the material is geared towards IoT (Internet of Things) devices, many of the recommendations

³⁵ NIST SP 800-82r3. Guide to Operational Technology (OT) Security. April 2022.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.ipd.pdf>.

³⁶ NIST SP 800-53 Revision 5. Security and Privacy Controls for Information Systems and Organizations. September 2020.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf>.

³⁷ NIST SP 800-213A. IoT Device Cybersecurity Guidance for the Federal Government: Establishing IoT Device Cybersecurity Requirements. November 2021.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-213.pdf>.

³⁸ NIST SP 800-213A. IoT Device Cybersecurity Guidance for the Federal Government: IoT Device Cybersecurity Requirement Catalog. November 2021.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-213A.pdf>.

and guidance put forth in the material can apply directly to managing the cybersecurity of operational technology (OT).

Cyberinfrastructure & Infrastructure Security Agency (CISA)

The Cyberinfrastructure & Infrastructure Security Agency (CISA)³⁹ provides a large collection of abstracts and links to source documents for existing recommended cybersecurity practices for a wide variety of control systems topics⁴⁰ such as patch management.⁴¹ Resources in this library have been developed and vetted by control systems subject matter experts (SMEs). Although many instances of OT deployments in NSF Major Facilities are highly customized or purpose-built per research program activity, the guidance and resources contained therein is an exceptional starting point to understand and implement cybersecurity practices in OT environments.

Industrial Control Systems Joint Working Group (ICSJWG)

The Cybersecurity and Infrastructure Security Agency (CISA) hosts the Industrial Control Systems Joint Working Group (ICSJWG)⁴² to facilitate information sharing and reduce the risk to the nation's industrial control systems. The goal of the ICSJWG is to continue and enhance the collaborative efforts of the industrial control systems stakeholder community in securing CI by accelerating the design, development, and deployment of secure industrial control systems.

MITRE ATT&CK® for ICS

The MITRE corporation produces a knowledge base called ATT&CK⁴³ to enumerate attacker tactics and techniques and “is used as a foundation for the development of specific threat models and methodologies.” MITRE has produced an ATT&CK Matrix tailored to ICS to help understand attacks upon operational technology and develop threat models for OT infrastructure.⁴⁴

³⁹ Cyberinfrastructure & Infrastructure Security Agency (CISA). U.S. Computer Emergency Readiness Team. <https://www.cisa.gov/uscert>.

⁴⁰ Cyberinfrastructure & Infrastructure Security Agency (CISA). Recommended Practices. <https://www.cisa.gov/uscert/ics/Recommended-Practices>.

⁴¹ U.S. Department of Homeland Security. Recommended Practice for Patch Management of Control Systems. December 2008. https://www.cisa.gov/uscert/sites/default/files/recommended_practices/RP_Patch_Management_S508_C.pdf.

⁴² Industrial Control Systems Joint Working Group (ICSJWG). <https://www.cisa.gov/uscert/ics/Industrial-Control-Systems-Joint-Working-Group-ICSJWG>.

⁴³ MITRE ATT&CK. <https://attack.mitre.org/>.

⁴⁴ MITRE ATT&CK for ICS. <https://attack.mitre.org/versions/v11/matrices/ics/>.

CIS Critical Security Controls ICS Companion Guide

The CIS controls represent a well prioritized set of cybersecurity controls which can be applied to secure infrastructure, including OT infrastructure. The “CIS Critical Security Controls ICS Companion Guide”⁴⁵ can be a useful resource for organizations attempting to develop secure configurations for OT assets. Note, a companion document does not yet exist for version 8 of the CIS controls.

NIST National Initiative for Cybersecurity Education (NICE)

Facilities and research programs can leverage the NIST National Initiative for Cybersecurity Education (NICE) framework⁴⁶ to orient themselves to the skills and knowledge necessary to operate secure environments. Tools and guidance offered therein is “aimed at helping employers, including human resource managers measure, assess, and build their cybersecurity workforces.” This resource can be used to readily identify personnel skills required to support OT, and can even be leveraged to easily produce full job descriptions for OT cybersecurity roles (e.g., the “PushButtonPD™ Tool”).⁴⁷

Research Security Operations Center (ResearchSOC)

The Research Security Operations Center (ResearchSOC)⁴⁸ at Indiana University helps make scientific computing resilient to cyberattacks and capable of supporting trustworthy, productive research. It does this by providing the operational cybersecurity services, training, and information sharing necessary to a community as unique and variable as research and education (R&E).

CI Compass, the NSF Cyberinfrastructure Center of Excellence (CI CoE)

CI Compass provides expertise and active support to cyberinfrastructure practitioners at NSF Major Facilities in order to accelerate the data lifecycle and ensure the integrity and effectiveness of the cyberinfrastructure upon which research and discovery depend.

7. Next Steps

It is the hope of the authors of this document that this roadmap will be useful to operators of operational technology in science, and NSF Major Facilities in particular, and also NSF

⁴⁵ CIS Critical Security Controls ICS Companion Guide.

<https://www.cisecurity.org/insights/white-papers/cis-controls-implementation-guide-for-industrial-control-systems>.

⁴⁶ National Initiative for Cybersecurity Education (NICE). NICE Framework Resource Center — History. <https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center/history>.

⁴⁷ National Initiative for Cybersecurity Education (NICE). NICE Framework Resource Center — Employer Resources. . <https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center/employer-resources>

⁴⁸ Research Security Operations Center (ResearchSOC). <https://researchsoc.iu.edu>.

Program Officers, so that they understand those gaps in securing OT in science and can better understand the need for prioritization and committing of resources to improving the state of securing OT in science.

At the same time, the intended audience for this document includes the Trusted CI organization itself, so it can best support NSF-sponsored organizations and scientific research facilities in securing their OT infrastructure and related programs.

NSF-sponsored organizations with scientific OT who would be interested in assistance from Trusted CI in discussing or implementing aspects of this roadmap is welcome to reach out to Trusted CI at any time: info@trustedci.org

Appendices

Appendix 1: The Trusted CI Framework's Foundation

Four Pillars, Sixteen Musts
A reasonable minimum standard for cybersecurity programs

Mission Alignment

1. Organizations must tailor their cybersecurity programs to the organization's **mission**.
2. Organizations must identify and account for cybersecurity **stakeholders and obligations**.
3. Organizations must establish and maintain documentation of **information assets**.
4. Organizations must establish and implement a **structure for classifying** information assets as it relates to the organization's mission.

Governance

5. Organizations must involve **leadership** in cybersecurity decision making.
6. Organizations must formalize roles and responsibilities for cybersecurity **risk acceptance**.
7. Organizations must establish a **lead role** with responsibility to advise and provide services to the organization on cybersecurity matters.
8. Organizations must ensure the cybersecurity program **extends to all entities** with access to, control over, or authority over information assets.
9. Organizations must develop, adopt, explain, follow, enforce, and revise cybersecurity **policy**.
10. Organizations must **evaluate and refine** their cybersecurity programs.

Resources

11. Organizations must devote **adequate resources** to mitigate cybersecurity risks deemed unacceptable by the organization.
12. Organizations must establish and maintain a cybersecurity **budget**.
13. Organizations must allocate **personnel** resources to cybersecurity.
14. Organizations must identify **external cybersecurity resources** to support the cybersecurity programs.

Controls

15. Organizations must adopt and use a **baseline control set**.
16. Organizations must select and deploy **additional and alternate controls** as warranted.

Appendix 2: Summary of Short-Term and Long-Term Recommendations

Term	Recommendations	[MF = NSF Major Facilities]
5.1 Mission (M)		
Short	<ol style="list-style-type: none">1. Account for OT as well as IT assets in MF strategic security plans2. Strategic plans and budgets must explicitly allocate resources for OT cybersecurity	
Long	<ol style="list-style-type: none">3. NSF should emphasize the centrality of OT cybersecurity in MF charters4. Work inside and across MFs to standardize IT and OT configurations5. Coordinate OT security across MFs, with Trusted CI and CI Compass	
5.2 Organization and Governance (OG)		
Short	<ol style="list-style-type: none">1. Appoint a cybersecurity lead with oversight over both IT and OT security2. Collaborate with host's existing IT and cybersecurity organizations3. Assess on-site OT expertise and awareness of cybersecurity topics4. Ensure the representation of OT personnel in institutional cybersecurity groups5. Develop guidelines for outsourcing expert cybersecurity roles	
Long	<ol style="list-style-type: none">6. Provision a Chief Information Security Officer or cybersecurity lead role7. Provision an IT/OT specialist and/or cybersecurity IT/OT expert role8. Creating a long term institutional goal for IT/OT collaboration and communication	
5.3 Policies (P)		
Short	<ol style="list-style-type: none">1. Include a review of OT infrastructure in regular assessment of assets2. Provide policy implementation and awareness training3. Require, through policy, that supplemental controls be applied to protect OT4. Establish a policy exception procedure to accommodate OT considerations	
Long	<ol style="list-style-type: none">5. Participate in industry forums to voice security concerns and support remediations	
5.4 Device Procurement and Maintenance (PM)		
Short	<ol style="list-style-type: none">1. Ensure the facility has a current inventory of OT environments and devices2. Educate decision-makers and purchasing entities on OT cybersecurity risks3. Implement technical controls to secure vendor and 3rd party access to OT devices4. Identify service-level activities that impact vendor support for OT equipment5. Add OT security requirements to vendor support contracts & research agreements6. Integrate cybersecurity requirements language into purchasing practices & policies	
Long	<ol style="list-style-type: none">7. Integrate cybersecurity supply chain risk management into policies/plans/practices8. Clearly communicate expectations and requirements to vendors9. Develop plans-of-action for vendor support limitations10. Develop and implement an OT acceptance testing program	
5.5 Technical Safeguards and Controls (T)		
Short	<ol style="list-style-type: none">1. OT systems should only run necessary tasks2. Document network and network-attached assets3. Identify network-attached control ports4. Apply the principle of least privilege to control access to OT assets5. Ensure the ability to recover from disaster6. Collaborating OT services should be segmented7. Control configurations of PCs and devices with external interfaces to OT networks	
Long	<ol style="list-style-type: none">8. Move to a Zero Trust Architecture	