



The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators

March 1, 2021

v1.0

Distribution: **Public**

Craig Jackson, Bob Cowles, Scott Russell,
Emily K. Adams, Ryan Kiser, Ranson Ricks,
Anurag Shankar

<https://www.trustedci.org/framework>
framework@trustedci.org



About Trusted CI

The mission of Trusted CI is to provide the National Science Foundation (NSF) community with a coherent understanding of cybersecurity, its importance to computational science, and what is needed to achieve and maintain an appropriate cybersecurity program.

Acknowledgments

Special thanks go to our Framework Advisory Board (FAB). The members are: Kay Avila (NCSA); Steve Barnett (IceCube); Tom Barton (University of Chicago); Jim Basney (NCSA); Jerry Brower (NOIRLab, Gemini Observatory); Jose Castilleja (NCAR/UCAR); Shafaq Chaudhry (UCF); Eric Cross (NSO); Carolyn Ellis (Purdue University); Terry Fleury (NCSA); Paul Howell (Internet2); Tim Hudson (NEON/Battelle/Arctic); David Kelsey (UKRI/WISE); Tolgay Kizilelma (UC Merced); Nick Multari (PNNL); Adam Slagell (ESnet); Susan Sons (IU CACR); Alex Withers (NCSA/XSEDE); Melissa Woo (Michigan State University). Thanks to our Framework Advisory Board Governance and Collaboration Leads: Andrew Adams and Von Welch.

Many thanks to current and past collaborators Kay Avila, Jim Marsteller, Kelli Shute, Susan Sons, Rebecca Yasky, John Zage, and the entire membership of the Large Facilities Security Team (<https://www.trustedci.org/lfst>). Their past work on this and related projects have had an important impact on the architecture and content of the Framework.

This document is a product of Trusted CI, the NSF Cybersecurity Center for Excellence. Trusted CI is supported by the National Science Foundation under grants numbered OCI-1234408, ACI-1547272, and ACI-1920430. For more information about Trusted CI, please visit <https://www.trustedci.org/>. Any opinions, 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, the FAB members, or their respective organizations.

Using & Citing this Work

This work is made available under the terms of the Creative Commons Attribution 4.0 International License. Please visit the following URL for details: <https://creativecommons.org/licenses/by/4.0/>.

Cite this work using the following information:

Jackson, Craig, Cowles, Bob, Russell, Scott, Adams, Emily K., Kiser, Ryan, Ricks, Ranson, and Shankar, Anurag, The Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators. 2021.

This work is available on the web at the following URL: <https://www.trustedci.org/framework>

Table of Contents

About the Trusted CI Framework	3
About this Framework Implementation Guide	5
Audience for the Guide	5
Community Engagement and Evolving the Guide	5
Structure of the Guide	6
Getting Started	7
Mission Alignment	8
Must 1: Mission Focus	9
Must 2: Stakeholders & Obligations	16
Must 3: Information Assets	20
Must 4: Asset Classification	26
Governance	31
Must 5: Leadership	32
Must 6: Risk Acceptance	36
Must 7: Cybersecurity Lead	43
Must 8: Comprehensive Application	48
Must 9: Policy	53
Must 10: Evaluation & Refinement	61
Resources	66
Must 11: Adequate Resources	67
Must 12: Budget	74
Must 13: Personnel	79
Must 14: External Resources	87
Controls	94
Must 15: Baseline Control Set	95
Must 16: Additional & Alternate Controls	101
Glossary of Key Terms	111
Appendix A - Does my Organization Need its own Cybersecurity Program?	114
Appendix B - Trigger Events	116
Appendix C: Baseline Controls Sets	117

About the Trusted CI Framework



The Trusted CI Framework is a tool to help organizations establish and refine their **cybersecurity programs**. In response to an abundance of guidance focused narrowly on cybersecurity controls, Trusted CI set out to develop a new framework that would empower organizations to confront cybersecurity from a mission-oriented, programmatic, and full organizational lifecycle perspective. Rather than rely solely on external guidance (which isn't tailored to the organization's mission and which may lack evidence of efficacy), the Trusted CI Framework recommends that organizations take control of their cybersecurity the same way they would any other important business concern: by adopting a programmatic approach. This framework is designed to be understandable and usable by non-cybersecurity and cybersecurity experts alike.

What is a cybersecurity program? A **cybersecurity program** is a group of related cybersecurity-focused projects and ongoing activities managed in a coordinated way to obtain benefits not available from managing them individually. Cybersecurity programs are an organ of the larger organization, living as part of that organization through its lifecycle.

Why take a programmatic approach to cybersecurity? Cybersecurity is a complex and dynamic space with evolving threats and technologies. Organizations cannot “solve cybersecurity” by adopting a checklist or rigidly adhering to a set of controls. By establishing a cybersecurity program, the organization acknowledges the complexity of the problem and commits to treating cybersecurity as an important, ongoing business priority.

Cybersecurity programs offer a number of benefits that cannot be achieved by simply implementing controls. Well-administered cybersecurity programs are:

1. **Focused on the organization's mission:** Cybersecurity programs are tailored to the needs, priorities, and risk tolerance of the organization and its mission.
2. **Ongoing and evolving:** Cybersecurity programs evolve with the organization as the organization matures. They include processes to adapt to changes in the organization's key assets, available resources, and place in the organizational lifecycle.
3. **About more than technology:** Cybersecurity programs address the full scope of cybersecurity decision making, including resourcing, governance, mission alignment, and control selection.

The Trusted CI Framework is structured around **4 Pillars** which make up the foundation of a competent cybersecurity program: **Mission Alignment, Governance, Resources, and Controls**. Composing these pillars are **16 Musts** that identify the concrete, critical requirements for establishing and running a competent cybersecurity program. The 4 Pillars and the 16 Musts combined make up the **Framework Core**, which is designed to be applicable in any environment and useful for any organization.

Visit www.trustedci.org/framework to learn more.

The Trusted CI Framework

Four Pillars. Sixteen Musts. An Architecture 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 they relate 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 address unacceptable cybersecurity risk.
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.

Visit www.trustedci.org/framework to learn more.

About this Framework Implementation Guide

A Framework Implementation Guide (FIG) is a document that provides detailed guidance on how to implement the Trusted CI Framework's 16 Musts for a specific audience or community. FIGs provide significantly greater detail than the Framework Core and include specialized guidance and recommendations for the target audience.¹

Audience for the Guide

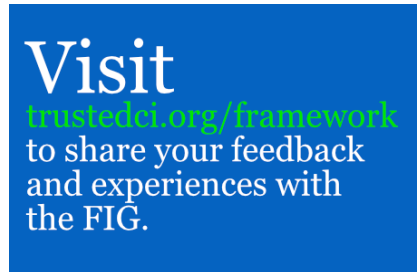
This Framework Implementation Guide (FIG) is designed for direct use by **research cyberinfrastructure operators (RCOs)**. We define RCOs as organizations that operate on-premises, cloud-based, or hybrid computational and data/information management systems, scientific instruments, visualization environments, networks, and/or other technologies that enable knowledge breakthroughs and discoveries.² These include, but are not limited to, major research facilities (*e.g.*, NSF Major Facilities, aka Large Facilities), research computing centers within research institutions, and major computational resources that support research computing.

While this definition is meant to be inclusive, it does not necessarily include organizations that do not operate their own computing infrastructure (*e.g.*, those that rely entirely on resources managed by a parent organization) or organizations that entirely outsource their IT and/or cybersecurity operations. For guidance on whether units of larger organizations, collaborations, or virtual organizations should establish their own cybersecurity programs, *see* **Appendix A**.

RCOs have a dual cybersecurity responsibility to protect their own information assets and also provide a cyberinfrastructure that eases the cybersecurity burden for research projects making use of that infrastructure.

Community Engagement and Evolving the Guide

As a product ultimately designed for use in the research and higher education communities, this Framework Implementation Guide was developed with significant input from stakeholders that represent a cross section of the target audience. This **Framework Advisory Board (FAB)** is a collection of 19 volunteers with diverse interests and roles in the research and education communities.³ From January 2020 through January 2021, Trusted CI's Framework project team engaged the FAB on a monthly basis, conducting 2 meetings per month to accommodate the broad



¹ Although FIGs are written with a particular audience or community's needs in mind, much of the guidance will prove valuable to readers outside of that audience or community.

² *See* <https://library.educause.edu/-/media/files/library/2009/4/epo0906-pdf.pdf>.

³ For a complete listing of the Framework Advisory Board and its members' respective positions, *see* the Acknowledgments section near the beginning of this FIG. Note, no NSF or other funding agency personnel participated as FAB members or authors of this Guide.

geographic distribution of the members. The FAB provided substantial input, suggestions, questions, and critiques during the drafting of the FIG content. Based on this input from the FAB, the authors refined and published v1.0.

Trusted CI will revise and refine this guide based on community feedback and experiences working with RCOs as they adopt the Framework and use its guidance.

Structure of the Guide

Organized by the four Pillars, each of the following chapters is focused on one of the sixteen **Musts**. Each chapter follows the same basic structure:

The Box. Each **Must** chapter begins with a box containing foundational content: The single sentence **Must** itself and a small amount of prose unpacking key definitions and the general meaning of the **Must**.


The BLUF. After the box, we briefly describe the relevance of this **Must** to RCOs and offer a short BLUF (bottom line up front) statement of our guidance on how RCOs should tackle the **Must**. This sets the stage for the remainder of the chapter.

Thereafter, each **Must** chapter has three common sections:

Why is this a Must? offers our rationale for identifying this as a requirement for standing up and maintaining a competent cybersecurity program. We describe both general reasons (that would apply to any organization) as well as particular motivations for RCOs to act.

The Roadmap offers a procedural view on how to get moving on the **Must**, as well as ongoing or repeated activities. For instance, the **Must 15 (Baseline Control Set)** Roadmap is broken into three steps: (1) Know your options; (2) Adopt the control set (or sets); and (3) Baseline the set.

Common Challenges & Recommendations provides brief discussions of common questions and blockers the RCO community faces and ideas about how to overcome them.

In **The Roadmap** and **Common Challenges & Recommendations** sections, we often make references to particular tools and templates that can make aligning to the **Must** more efficient. We highlight these with the hammer icon: 

Getting Started

Adopting the Trusted CI Framework does not require mindless implementation of all our guidance. It does require acknowledging that the 16 **Musts** need to be addressed to have a competent cybersecurity program and making a concerted effort to act on them all. The implementation guidance in this document is designed to give RCOs a reasonable path, the best tools we can develop or find, and ways to address common challenges.

This guide is a reference and a resource, and is not designed to be read cover to cover like a novel, or implemented in a strict sequence like a cake recipe. The following are some ideas about how to get started using the Framework and this guide.

Read the opening material for each chapter. The beginning of each chapter contains definitions and context for what the Must is all about. Unless we have a strong recommendation about a particular implementation path, most of the “how” details come later. Come back to the details of the chapters after you’ve made some decisions about priority.

Consider where your RCO is in its lifecycle. If your RCO is pre-operational and you are thinking seriously about cybersecurity, good for you! Newly founded or construction-phase RCOs have the opportunity to think strategically and holistically about how to make progress on the 16 Musts. We strongly encourage you to give attention to the first three Pillars: Mission Alignment, Governance, and Resources. If Controls are where the virtual “rubber meets the road” in cybersecurity, the first three Pillars help you build the drivetrain that makes the journey possible.

If your RCO is operational, consider some focused effort on **Must 10 (Evaluation & Refinement)** evaluation activities. And, if you don’t have a baseline control set selected (*see* **Must 15 (Baseline Control Set)**) or implemented, you may need to start moving on the most fundamental, basic, universally applicable controls. You may not only need to turn back the tide of low sophistication attacks, but detect ongoing events about which you are not yet aware.

Prioritize effort, but keep a broad view. Most RCOs, but certainly those with fledgling cybersecurity capabilities, will need to prioritize efforts to evaluate and act on the Musts. Based on your situation, you may not have the luxury of systematically pushing all 16 “rocks” up the “hill” to programmatic maturity. That said, we encourage you to avoid focusing so much on one or two Musts that you miss opportunities to capitalize on the interconnectedness of the Framework.

Find a way to progress on Mission Alignment. Whether your RCO is brand new or has been operational for decades, whether your cybersecurity capability is advanced or playing catch-up, the four Mission Alignment Musts should not be overlooked or sidelined. The first Pillar is the most fundamental.

Reach out. If you are struggling with how to prioritize and where to begin, reach out to your peer organizations and Trusted CI for perspectives and help.



Mission Alignment



Must 1: Mission Focus

Organizations must tailor their cybersecurity program to the organization's mission.

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.

A cybersecurity program exists to support its organization's mission.⁴ As such, the cybersecurity program must ensure that the security controls, processes, and structures it recommends are designed to maximize their benefit to the mission. After all, security is not without downsides: it can cost money, time, resources, functionality, and/or goodwill. But security also plays a powerful role in enabling the mission: ensuring reliable, consistent, and verifiable outcomes. Security plays both a mission-enabling and a risk-reducing function. Cybersecurity decision makers should tailor the cybersecurity program to maximize the benefits of security while minimizing the burdens. Moreover, security will impact different organizations differently, depending on a host of factors, including their specific mission, culture, and history: the cybersecurity program should be tailored to account for the organization's unique needs.

Effectively tailoring the cybersecurity program requires an understanding of the organization's mission. Although missions are often referred to in the singular (*i.e.*, "the mission"), this usage appears to signify a generalized alignment with high-level organizational goals, (*e.g.*, "in support of the mission"), not that there exists only one mission. Indeed, missions exist at all levels of organizational complexity, from as simple as a pencil to as complex as an entire organization or nation, and these different missions interrelate. The cybersecurity program must be tailored to reflect 1) the organization's stated mission (*i.e.*, in the organization's mission statement); 2) the missions the organization operates in service of; and 3) the missions that directly support the broader mission. For instance, an organization's cybersecurity program should consider both the mission(s) of any parent organization and the missions of the projects and programs critical to the mission's success. Therefore, organizational missions must take into account and prioritize many overlapping concepts, such as corporate values, culture, goals, and fears.

Organizational missions of research cyberinfrastructure operators (RCOs) often involve supporting research, researchers, and science. However, RCOs may have broader and more aspirational missions as well, such as benefiting society or the planet. An RCO's lower-level missions may include supporting one-of-a-kind scientific instruments, collecting or processing critical research data, or providing a platform for the public to access research. Yet RCOs are not immune from the missions

⁴ The definition of mission used here was derived from a combination of business, military, and academic sources. This definition is intended to address both the complexity of the mission concept and the reality of how organizations make decisions. Our definition intends to capture both what organizations say they care about and what their actions suggest they actually care about. See https://www.doctrine.af.mil/Portals/61/documents/Volume_1/V1-D34-Levels-of-War.pdf. *Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0*

more commonly associated with private businesses, and must also consider the roles profit, reputation, and competitive advantage play in their organizational decision making.⁵

Finally, it is important to note that many RCOs are embedded within or affiliated with a larger parent institution, such as a research university, and that their mission will often need to take into account the mission of that parent institution.

Why is this a Must?

“If the highest aim of a captain were to preserve his ship, he would keep it in port forever.” - Thomas Aquinas

Cybersecurity is not undertaken as an end unto itself: its purpose is to support the organization’s mission. Cybersecurity is a protective, risk-reducing domain that is about supporting an environment of deliberative risk-taking.

Tailoring cybersecurity is crucial because security, in and of itself, is not the mission, nor does it inherently advance the mission: If your mission is to manufacture shoes, cybersecurity does not directly produce any shoes. Rather, cybersecurity is a response to risks and threats in the environment that, when left unmitigated, could impede shoe production. Therefore, its value must always be understood as relative and contingent upon the risks and threats faced.⁶ (Note, however, that many security activities can have significant positive impacts on other aspects of the organization, such as quality assurance and business continuity.)

A poorly tailored cybersecurity program is a bad investment for the organization. Overinvesting in cybersecurity results in increasingly diminishing returns, wasting resources that could have been used to directly advance the mission. Similarly, underinvesting in cybersecurity is also a bad investment, as an investment in preventative cybersecurity can be significantly smaller than the cost of a cyber incident,⁷ and increasingly contracts and partnerships are including cybersecurity requirements as a condition. Moreover, cybersecurity needs can vary greatly among different sectors and organizations. Programs should be tailored to meet the specific needs of the organization or else cybersecurity won’t be maximizing its value to the mission.

The Roadmap

This section describes the steps needed for an RCO to tailor their cybersecurity program to the organization’s mission. Doing so is an ongoing process that must take into account the present and future needs and priorities of the RCO. New and/or immature RCOs can have significantly different requirements from older and more mature RCOs. The cybersecurity program’s leadership should be

⁵ Other common mission factors to consider include third party dependencies and externalities, safety, legal compliance, and ethics.

⁶ It is important to emphasize that organizations can never completely eliminate risk. Instead, the program should help the organization decide how much cyber-related risk it is willing to accept and address risks it deems unacceptable. *See* also **Must 6 (Risk Acceptance)**. There is no correct amount of risk, and reasonable organizations can differ on how much risk they are comfortable with accepting. The cybersecurity program should be tailored to match these preferences.

⁷ *See, e.g.,* <https://www.ibm.com/security/data-breach>.

evaluating how best to enable its RCO's mission on an ongoing or periodic basis. *See* **Must 10 (Evaluation & Refinement)**. Although the process is ongoing, a useful starting point is to think of it as four basic steps: Step 1 is to understand the mission; Step 2 is to create a cybersecurity program strategic plan; Step 3 is to implement the plan; and Step 4 is to evaluate and adjust the plan and understanding of the mission.

Note: The process of evaluating and tailoring the organization's cybersecurity program should be present throughout all of the Musts. When implementing each Must, leadership must consider the benefits and burdens that security imposes, and strike an appropriate balance.

→ Step 1. Understand the Mission.

Although seemingly straightforward, the “mission” is a complex concept. Organizations often have a partial, incomplete, or misguided understanding of their missions. The RCO's mission may be much more complex than the RCO's mission statement. (*See* discussion below in Common Challenges). Moreover, the security program's design and implementation need to account for not only the highest-level organizational mission, but also the lower-level missions of, for instance, individual projects, programs, and initiatives that the organization undertakes or depends upon.

Understanding the organization's mission requires a discovery process, and involves engaging relevant stakeholders to understand and account for the organization's mission. *See* **Must 2 (Stakeholders & Obligations)** for the range of stakeholders with an interest in the cybersecurity program, and **Must 5 (Leadership)** for more discussion of how to engage RCO leadership.

A useful framing device is to decompose “the mission” into three nested categories:⁸ 1) the highest-level strategic or organizational missions; 2) the middle-level ongoing operational or programmatic missions which advance the strategic missions; and 3) the lowest-level tactical or project-level missions in support of an operational mission.

Organizational missions operate at the highest level of abstraction, comprising the organization's goals, motivations, and priorities writ large. An example of an RCO's organizational mission might be to advance science by providing low-cost, long-term data storage and curation for researchers. (RCOs will typically have a small number of organizational missions.)

Programmatic missions operate at a middle level of abstraction, representing ongoing activities with clear objectives but without near-term end-points. An example of an RCO programmatic mission might be to maintain supercomputer operations.

Finally, project-level missions are discrete, time-limited, and project or event-driven activities with clear objectives and endpoints. An example of an RCO project-level mission would be to install a major software update to its research services with low downtime. (Project-level missions may be numerous enough that enumerating them is infeasible or not useful.)

⁸ These categories derive from the strategic, operational, and tactical levels of war. *See, e.g.*, “Levels of War, Air Force Doctrine,” https://www.doctrine.af.mil/Portals/61/documents/Volume_1/V1-D34-Levels-of-War.pdf.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ Step 2. Develop a Cybersecurity Program “Strategic Plan”.

The second step is to develop a strategic plan for tailoring the cybersecurity program. The goal here is to document the RCO’s overarching strategy explaining how the program will be tailored to the organization’s mission (*i.e.*, who, what, when, where, why, and how). The specifics of the strategic plan will vary based on the needs of the RCO, but strategic plans should generally include: 1) a mission statement for the cybersecurity program, 2) a statement of the program’s cybersecurity strategy, and 3) a timeline identifying key programmatic milestones. Trusted CI offers a **Cybersecurity Program Strategic Plan Template**⁹ to aid RCOs in this process.

The first component is the **cybersecurity program’s mission statement**.¹⁰ This should state the purpose of the cybersecurity program, and how it supports the mission of the organization (including any relevant parent or third party missions). This statement serves as a communication tool within the cybersecurity program and will guide decision making.

The second component is a **cybersecurity strategy** for the program. This strategy is a brief statement outlining the priorities and direction of the program and explaining how it advances the RCO’s mission. The cybersecurity strategy should focus on higher-order security considerations¹¹ and should directly reference the RCO’s missions. For instance, one RCO’s cybersecurity strategy may prioritize business continuity and rapid recovery from incidents (*i.e.*, Fault Tolerance), whereas another may prioritize ensuring data integrity. A cybersecurity strategy can also identify a desired end state: *i.e.*, what the program should look like at the end of the stated timeline.

The third component is a **timeline** outlining key programmatic **milestones** the RCO seeks to achieve during the relevant time period. This timeline should focus on major developments for individual Musts, such as adopting a baseline control set for **Must 15 (Baseline Control Set)** or hiring a cybersecurity lead for **Must 7 (Cybersecurity Lead)**. Additionally, the timeline should include a brief explanation about *why* these milestones were prioritized and how they advance the RCO’s mission. For instance, for an RCO that is first establishing their cybersecurity program, they may entirely prioritize **Must 15 (Baseline Control Set)** for the first year, stating that they believe that adopting and utilizing a baseline control set will provide the greatest overall value to the RCO’s mission.

In addition to these three components, RCOs can provide more detail to the strategic plan as they see fit. For instance, an RCO could walk through each **Must** briefly describing how to tailor the **Must** to the RCO’s mission, including writing down the mission trade-offs and assumptions that were considered and explaining how the documented course will provide value to the RCO.

Finally, when drafting (and implementing) the strategic plan, it is valuable to consider: 1) taking advantage of any parent organization’s cybersecurity programs when available; 2) connecting with

⁹ Check out Trusted CI’s **Cybersecurity Program Strategic Plan Template**, available at <https://trustedci.org/framework>.

¹⁰ <https://modernciso.com/2017/08/22/cyber-resilience-a-primer-part-1-defining-your-security-programs-mission-statement/>.

¹¹ Potential resources for describing high level security concerns include the Information Security Practice Principles <https://cacr.iu.edu/about/principles.html> and the CIA triad <https://resources.infosecinstitute.com/cia-triad/>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

and using lessons learned from peer organizations; and 3) seeking out expert assistance when tackling unique or particularly challenging problems.

→ Step 3. Approve and Implement the Strategic Plan.

The third step is to gain leadership approval for the strategic plan and begin implementing the actions and priorities it sets out. The process of gaining leadership approval is invaluable, as it provides an opportunity for RCO leadership and the cybersecurity lead to have a discussion and reach consensus on the long term trajectory of the program.

Once approved, Step 3 is where resources are allocated to implement the plan and where detailed project plans, timelines, and schedules are drafted. Since the plan created in Step 2 is strategic, it will be primarily focused on high-level goals and priorities, and will require more work to detail the operational and tactical decisions that will help achieve those goals. To aid in these decisions, consult with the corresponding FIG chapters for each Must, in particular the Roadmap and Common Challenges and Recommendations sections. Implementing **Must 1** is an ongoing activity and will be implicated in the Roadmaps for each other Must.

→ Step 4. Evaluate & Adjust.

The final step involves evaluating the success of the current tailoring strategy and making adjustments where appropriate. Tailoring an organization's cybersecurity program is not a "one and done" initiative: it is something the organization will need to periodically do to ensure the program continues to enable the organization's mission. Missions evolve and adapt over time, risks and threats in the environment change or are better understood, stakeholders update their requirements, and the organization may simply learn from the experiences of implementing its security strategy. For instance, organizations may experience a change in stakeholder obligations that requires a change in the prioritization of Musts or the overarching strategy initially set out in the strategic plan.

Evaluating and adapting the cybersecurity program's strategic plan should be done periodically. If adjustments are made too frequently, an organization will not have time for Must implementations to progress, and the assessments will be premature. While if done too infrequently, the organization may be stuck with a poorly-tailored cybersecurity strategy that inhibits the mission. A good rule of thumb is to engage in this process every three to five years. For a discussion of evaluation and refinement of the cybersecurity program rather than the strategic plan, *see* **Must 10 (Evaluation & Refinement)**.

RCOs should consider the following tool:



Trusted CI maintains a number of templates to aid organizations in establishing their cybersecurity program.¹² In particular, the **Cybersecurity Program Strategic Plan Template** provides a good starting point for documenting the cybersecurity program's mission statement, cybersecurity strategy, and key milestones.

¹² <https://trustedci.org/framework>.

Common Challenges & Recommendations

This section describes some common **Must 1** challenges and offers recommendations on how to overcome them.

→ Organizations don't know their mission.

Organizations may have a vague or partial understanding of their mission. Organizations need a structured way of considering what they do, what they care about, and what their priorities are.¹³ Although no single, definitive process can be used for determining an organization's mission, some of the following options may prove useful.

RCOs may consider conducting interviews or structured discussions with leadership, stakeholders, and project-level management.¹⁴ RCOs can also explore hypothetical scenarios pitting competing RCO interests against each other. Or RCO leadership can simply ask questions such as “what do we do,” “who do we serve,” and “what makes us different.” The goal is to explore and document what the RCO cares about, and what factors are most salient in driving RCO decision making. It may prove helpful to engage a disinterested third party to provide a neutral perspective. Regardless of which method is used, the final product should be vetted by the RCO's leadership to ensure that they agree with its statement of priorities.

A common output is the “mission statement,” but depending on how they are constructed, these are often of limited value, stating vague goals, relying entirely on marketing language, or simply reflecting an unrealistic ideal of the organization's mission.¹⁵

Note, a pitfall organizations face when thinking critically about their mission is to become unnecessarily specific, outlining not only organizational objectives but also *how* those objectives should be achieved.¹⁶ This pitfall unnecessarily restricts the organization's options and obfuscates the true mission. When considering their mission, organizations should always ask themselves, “Is this something I care about intrinsically, or do I think this is necessary to achieve something I care about?” In the latter case, the organization should resist considering that as part of its mission.

→ Need to Act Quickly?

Ideally RCO decision makers will have sufficient time to deliberate and think critically about how to

¹³ Note, the process of developing a mission statement specifically for the cybersecurity program (*see* Step 2) can also benefit the organization's more generalized understanding of its missions, and may even prompt a reevaluation of the organization's mission statement. *See* <https://modernciso.com/2017/08/22/cyber-resilience-a-primer-part-1-defining-your-security-programs-mission-statement/>.

¹⁴ Note, these options for understanding the organization's mission can be conducted by a wide range of actors. Although the cybersecurity lead is the natural starting place, these responsibilities may be delegated, may be managed by a committee, or may be undertaken by the RCO's leadership.

¹⁵ *See, e.g.*, <https://www.nytimes.com/2007/09/23/jobs/23mgmt.html>.

¹⁶ Note, in the particular case of organizational mission statements, organizations may have different motivations regarding what details to include. Mission statements often serve as external communication tools or marketing products, and the details that are included are likely to reflect this fact. As such, the details included in the mission statement should not necessarily be viewed as coextensive with the true organizational mission.

best tailor its cybersecurity program. However, this may not always be the case. Existing organizations may need to establish and tailor their cybersecurity program “in a hurry” and want a way to get things up and running quickly. The simplest approach for quickly tackling **Must 1** would be to draft a working mission statement and outline provisional strategies to get each Must to “good enough.” Although the goal is ultimately to tailor each Must to the RCO’s mission, the first priority in this scenario is to get each Must up and running in at least a minimal capacity. This may shift the analysis from “what is best for us?” to “what can we do right now?” The RCO’s leadership may also determine that some Musts are more important than others to get up and running quickly, and prioritize those Musts. From there leadership can begin selectively improving Musts based on the anticipated benefit to the RCO’s mission.

→ Explaining “Why?”

One of cybersecurity professionals’ primary roles is to communicate *why* cybersecurity is good for the mission. For example, explaining how a seemingly burdensome control can help the organization avoid worse problems down the road. To do so, cybersecurity professionals must be aware of security’s burdens, and be able to justify that burden in terms of benefit to the organization’s mission. Investing in security costs time, money, and human capital, and those are resources that the organization is not investing in other endeavors. Security is also potentially disruptive, placing hurdles between the organization’s personnel and their tasks (*e.g.*, multiple authentication requirements; complex passwords; restricted access to tools and resources). And perhaps even more challenging, specific security controls may prove unpopular within the organization, particularly when the value of those controls is not well communicated.

However, often the most impactful cybersecurity communicators are organizational leaders themselves. Leaders are particularly well-situated to communicate to the entire organization, and to ensure that the communications are listened to. To maximize the effectiveness of cybersecurity initiatives, leaders should take a forward role in communicating about cybersecurity to the organization.

To communicate effectively, cybersecurity should be contextualized in terms of the mission by both cybersecurity professionals and by leadership. Their communications should acknowledge the burdens, contextualize them in terms of corresponding benefit, and explain why they are appropriately tailored to maximize the benefit to the mission.



Must 2: Stakeholders & Obligations

Organizations must identify and account for cybersecurity stakeholders and obligations.

Cybersecurity stakeholders are people or entities with interest in or affected by an organization's cybersecurity. Cybersecurity obligations are any internally or externally imposed processes or practices that impact the operation of the organization's cybersecurity program. Accounting for these stakeholders and obligations involves making and following through on conscious, documented decisions with regard to them.

The cybersecurity program of a research cyberinfrastructure operator (RCO) involves both internal and external stakeholders. Internal stakeholders include the RCO and IT leadership, application developers, system administrators, and information system users. External stakeholders include research projects, suppliers, parent organizations, sponsors, and consortia for many activities (*e.g.*, federated identity management, incident response, threat monitoring). Human Resources, legal counsel, and even the cybersecurity team may be internal or external to the RCO depending on the size of the RCO and its relationship to any parent organization. Researchers are an important class of stakeholders for an RCO: they are critical to the RCO's mission as their work is a primary reason for the existence of the RCO. Researchers supported by the RCO may include internal researchers on staff and external researchers collaborating on projects using the RCO's cyberinfrastructure.

Flowing from stakeholders, cybersecurity obligations are any internally or externally imposed processes or practices that impact the operation of the organization's cybersecurity program. In addition to guidance like The NSF Major Facilities Guide¹⁷ or the terms and conditions from a funding agency,¹⁸ RCOs may be subject to specific data use agreements or statutory, regulatory, contractual, or other legal requirements.¹⁹ These obligations set expectations for the cybersecurity control implementation or programmatic.²⁰ Regulations may be international, national, state, or local and might be specific to a particular type of information asset (*e.g.*, information privacy requirements of HIPAA,²¹ FERPA,²² and GDPR²³ and export control restrictions in ITAR²⁴ and EAR²⁵).

Why is this a Must?

Fulfilling obligations and serving stakeholder interests may be central to an RCO's mission, and also

¹⁷ https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf19068.

¹⁸ https://www.nsf.gov/awards/managing/co-op_conditions.jsp.

¹⁹ Federal or state privacy laws, contractual obligations to apply security controls such as NIST 800-171 for CUI.

²⁰ <https://blog.trustedci.org/2017/06/nist-sp-800-171-and-its-potential.html>.

²¹ <https://www.hhs.gov/hipaa/index.html>.

²² <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>.

²³ <https://gdpr-info.eu/>.

²⁴ https://www.pmdotc.state.gov/ddtc_public?id=ddtc_kb_article_page&sys_id=%2024d528fddbfc930044f9ff621f961987.

²⁵ <https://www.bis.doc.gov/index.php/regulations/export-administration-regulations-ear>.

a source of complexity, burden, and tough decisions about priorities. Both stakeholders and obligations can define, change, promote, or constrain organizational missions and activities. A prerequisite to effective mission-oriented decision making is understanding who may be interested or affected by, or have expectations and obligations on, the organization's cybersecurity program as well as applicable standards, rules, or other requirements. In many, but not all cases, there is a trust relationship with the stakeholder.

The failure to adequately account for cybersecurity requirements may carry the threat of violating a data use agreement or regulatory enforcement, fines, and other punitive actions taken against the organization for failure to comply. Privacy laws like HIPAA, FERPA, and GDPR carry potentially significant financial penalties for failure to comply, whereas contractual requirements like those imposed for handling Controlled Unclassified Information (CUI) have. In contrast, the contracts have the threat of breach of contract lawsuits, litigation under the False Claims Act,²⁶ or merely the failure to bid on future contracts successfully.

The Roadmap

This section describes the steps needed for an RCO to identify and account for cybersecurity stakeholders and obligations. Organizations should involve various stakeholders (*e.g.*, legal counsel, human resources, grant, and contract personnel) to aid in identifying external requirements for cybersecurity and privacy that have an impact on the program. Many of the previously mentioned stakeholders can also be a source of information and help. Note that the requirements are often not prescriptive and leave the opportunity for interpretation and alternative methods of satisfying them.

→ Step 1. Stakeholder Identification.

Who are the RCO's stakeholders? As mentioned, cybersecurity stakeholders are people or entities with an interest in or affected by the organization's cybersecurity. Cybersecurity obligations are any imposed processes or practices that impact the operation of the organization's cybersecurity program.²⁷ Given this definition, the RCO should account for the entities that meet the parameters of this definition by understanding their interest, involvement, interdependencies, influence, and impact on the cybersecurity program's effectiveness. Identifying stakeholders enables the RCO to know the appropriate engagement focus for each stakeholder or group of stakeholders.²⁸

→ Step 2. Opportunities and Obligations Identification.

The stakeholders' effect on the cybersecurity program may be opportunities benefiting the RCO or restrictions affecting the program itself (*e.g.*, lack of organizational support, funding) or a source of obligations (*e.g.*, involving protected data). The obligations “... may be of a political, cultural, or strategic nature; they may be territorial, organizational, structural, functional, personnel, budgetary,

²⁶ <https://www.csoonline.com/article/3518728/recent-false-claims-act-cases-a-caution-to-govt-contractors-that-skimp-on-security.html>.

²⁷ Definition adopted from the Project Management Institute (PMI) Project Management Body of Knowledge (PMBOK) Guide, 6th Edition.

²⁸ Project Management Institute (PMI) Project Management Body of Knowledge (PMBOK) Guide, 6th Edition.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

technical, or environmental ...”²⁹ A spreadsheet of obligations and stakeholders is a good way to keep track as they are identified. At times, new obligations arise. Guidance for evaluating and refining the cybersecurity program to cover those obligations can be found in **Must 8 (Comprehensive Application)** and **Must 10 (Evaluation & Refinement)**.

→ **Step 3. Stakeholder Engagement Strategies.**

Equally important to identifying stakeholders and obligations is developing an engagement strategy for each. RCOs need to understand and manage the expectations of stakeholders through timely planning of interactions to produce the best impact on their program.³⁰ This planning is particularly important when cybersecurity-related activities are dependent on stakeholder decisions. RCO should document stakeholders and prepare an appropriate engagement approach for each.^{31,32}

It may be useful to think about each stakeholder by placing them in one of four quadrants.³³

<p>High Power / Low Interest <i>Keep Satisfied</i></p> <p><i>“Those with high power but low interest should be kept satisfied and ideally brought around as patrons or supporters for the proposed policy change.”</i></p> <p>It is important to keep them informed of significant changes but not inundated them about lesser topics.</p>	<p>High Power / High Interest <i>Engage Closely and Influence Activity</i></p> <p><i>“Stakeholders with high power, and interests aligned with the project, are the people or organisations it is important to fully engage and bring on board. If trying to create policy change, these people are the targets of any campaign.”</i></p>
<p>Low Power / Low Interest <i>Monitor (minimum effort)</i></p> <p>Advise them of significant changes and monitor for movement to one of the other quadrants.</p>	<p>Low Power / High Interest <i>Keep Informed</i></p> <p><i>“Stakeholders with high interest but low power need to be kept informed but, if organised, they may form the basis of an interest group or coalition which can lobby for change.”</i></p>

“‘Interest’ measures to what degree they are likely to be affected by the research project or policy change, and what degree of interest or concern they have in or about it. ‘Power’ measures the influence they have over the project or policy, and to what degree they can help achieve, or block, the desired change.”³⁴

²⁹ Computer and Information Security Handbook, 2nd edition (2013), Sokratis K. Katsikas, Chapter 53, Risk Acceptance, pg 910.

³⁰ “Plans are worthless, but planning is everything” Dwight D Eisenhower, <https://babel.hathitrust.org/cgi/pt?id=miaa.4728417.1957.001&view=1up&seq=858>.

³¹ Project Management Institute (PMI) Project Management Body of Knowledge (PMBOK) Guide, 6th Edition.

³² In some cases, the RCO may need to conduct engagements through a vendor representative or leadership channels. When developing stakeholder engagement strategies, consider this when direct engagement is not possible.

³³ <https://www.odi.org/sites/odi.org.uk/files/odi-assets/publications-opinion-files/192.pdf>.

³⁴ <https://www.odi.org/sites/odi.org.uk/files/odi-assets/publications-opinion-files/6459.pdf>.

Develop stakeholder communication plans according to power and interest, and an engagement strategy for each classification. As a bonus, produce for key stakeholders a description of the specific ways the cybersecurity program enables the mission and stakeholders associated with each. In cases where there is no trust relationship with the stakeholder, consider building one.

→ Step 4. Implement.

The actual cybersecurity program implementation of this Must occurs in other Musts within the Framework. See **Must 9 (Policy)**, **Must 10 (Evaluation & Refinement)**, **Must 14 (External Resources)**, **Must 15 (Baseline Control Set)**, and **Must 16 (Additional & Alternate Controls)**.

Common Challenges & Recommendations

This section describes some common **Must 2** challenges and offers recommendations on how to overcome them.

→ Tracking Adherence.

The cybersecurity program should include procedures and accountability measures to track the RCO's status with respect to the requirements and obligations to cybersecurity stakeholders. Appropriate accountability measures will vary based on specific stakeholder requirements and can include, but are not limited to, assessments, incident tracking metrics, risk response metrics, and cybersecurity control efficacy. In some cases, such as with protected health information, the methods for tracking an RCO's status will be defined according to regulatory requirements and efforts will by necessity be focused on enabling an RCO to adhere to these requirements. In other cases stakeholders such as funding agencies may have requirements such as reporting requirements which inform the methods of tracking the RCO's status.

→ Ensure Open Communication with Important Stakeholders.

Cybersecurity personnel may feel reluctant to engage with stakeholders (even within the RCO or parent organization) for fear of inviting conflicts, complexity, and churn. Stakeholders, including other RCO personnel, may have the same fears about engaging cybersecurity.³⁵ However, there can be serious legal or financial consequences for when stakeholders, including departments like human resources, procurement, and legal counsel, are out of alignment. Such service departments have a similar role to cybersecurity in protecting the organization from harm. We recommend that RCO leadership and cybersecurity personnel set an expectation of stakeholder awareness and communication that emphasizes relationship building and situational awareness before “things go boom.”

³⁵ This excuse is sometimes given for why researchers resist communicating with cybersecurity personnel. <https://www.forbes.com/sites/forbestechcouncil/2020/06/12/cyber-defenders-rebranding-cybersecurity-for-non-techies/#3e5e61cc13b9>.



Must 3: Information Assets

Organizations must establish and maintain documentation of information assets.

Information assets are valuable, sensitive, and/or mission critical information³⁶ and information systems.³⁷ Information asset documentation is the collection of artifacts describing the cybersecurity relevant details of information assets presented in a form that is useful to cybersecurity professionals and decision makers.

For cybersecurity purposes, the documentation may be a collection of artifacts that represent the *what* and *who* the cybersecurity program protects, where they are located, and how they are connected. In the aggregate, the documentation may consist of diagrams, automated and/or manually-supplied data in a variety of formats and files.

Documentation should contain the cybersecurity relevant details needed to facilitate organizational visibility and oversight for the asset. The core of this documentation is typically an inventory of assets which includes the key details about the assets. These details should include, at a minimum, an identifier, the asset's location, an asset owner,³⁸ and information asset classification in relation to the organization's information asset classification structure. *See Must 4 (Asset Classification)* for further discussion of asset classification. Additional details should be included where useful. These may include the asset's version, asset tag ID, warranty/support and expiration, interactions with other assets, or missions the asset supports. The decision regarding the details to include will vary depending on the asset and on the organization, with decision makers determining what information is necessary.

Documentation can take on many forms and will vary significantly based on the specifics of the asset and the maturity of the organization. For some organizations, their documentation may be limited to a manually populated excel spreadsheet listing assets over a certain dollar amount; whereas for other organizations, their documentation may provide dynamically-generated, in depth details. Well maintained documentation of information, systems, and components should be considered a basic necessity for a successful cybersecurity program.

³⁶ **Information** is any communication or representation of knowledge such as facts, data, or opinions in any medium or form, including textual, numerical, graphic, cartographic, narrative, or audiovisual. Organizational information may be stored and used within the organization's information systems, as well as flow out to third party systems. *See* National Information Assurance (IA) Glossary, CNSS Instruction No. 4009, Apr. 2010.

³⁷ An **information system** is a discrete set of information and related resources (such as people, equipment, and information technology) organized for the collection, processing, maintenance, use, sharing, dissemination, and/or disposition of information. These include, but are not limited to, mobile devices, routers, servers (the usual commodity IT equipment), as well as industrial control systems (ICS) / supervisory control and data acquisition (SCADA) systems, physical security systems, heating, ventilation, and air conditioning (HVAC) systems, and any other connected devices. *See* 44 U.S.C. 3502.

³⁸ **Must 6 (Risk Acceptance)** and **Must 8 (Comprehensive Application)** contain additional discussion and guidance regarding identification and assignment of appropriate asset owners.

Why is this a Must?

Accurate information asset documentation is one of the most powerful tools available to the information security team. Every other step in the process of creating a healthy cybersecurity program begins with knowing what information and information systems are to be secured, who is responsible for the assets, and how they are interdependent. Documentation of an organization's information assets is a prerequisite for 1) the assignment of decision making authority regarding those assets; 2) asset classification activities; 3) control selection and implementation; and 4) effective incident response. Furthermore, asset documentation is a necessary component for an organization to maintain the capability to perform these activities over time. These activities in turn will provide opportunities to further improve documentation. Incident response and asset classification activities for example will often uncover additional information about an environment which will be important to capture in documentation, which will then in turn be valuable for performance of those activities in the future. Indeed, this is one of the reasons why almost every widely used baseline control set includes documentation requirements (*e.g.*, asset inventories) as one of their controls.³⁹

Without effective documentation, organizations will suffer from a host of problems that arise from poor visibility into and understanding of one's own systems. Assets will be neglected because they aren't systematically tracked and assigned a person to oversee them. Institutional knowledge will be lost when key personnel leave because of the difficulty in educating their replacements on the environment. Neglected assets provide additional targets for attackers as software ages and vulnerabilities arise. Decision makers will be unaware of the gaps in protection they have because segments of their organization are invisible to them.

The Roadmap

This section describes the steps needed for a research cyberinfrastructure operator (RCO) to develop information asset documentation. Step 1 is where the organization determines what asset documentation it needs for its mission. Step 2 is about using the tools at one's disposal to gain visibility into the organization. Step 3 is about creating documentation that is valuable to the organization's cybersecurity professionals and decision makers.

→ Step 1. Determine need.

Asset documentation can be varied and distributed, and RCOs will have significant discretion in how they choose to document their assets. As such, a preliminary step for producing asset documentation is to determine what types of documentation the organizational missions would benefit from, what information would ideally be included in that documentation, how that information should be presented, and the resources available to produce and maintain the documentation. Documentation

³⁹ Several of these control sets go further by separating these documentation activities into distinct controls for a variety of reasons. Some such as the CIS Controls Version 7.1 separate software inventory and hardware inventory into distinct controls due to the different activities required to effectively develop and maintain documentation of each. NIST 800-53a groups controls into related "Control Families", most of which include documentation activities relevant to the topic addressed by the control family. For example, RA-2 specifies that organizations document the categorization of different assets while CM-8 describes how an organization should document system components.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

can take a number of forms and can exist at multiple levels of abstraction, so organizations need to think critically about where it requires the greatest visibility.

The primary area of first concern for most organizations will be identifying and cataloguing their critical assets as well as establishing greater visibility into those mission-critical assets and their key dependencies. To accomplish this, organizations should start by cataloguing information themselves rather than information systems. The mobility and easy reproducibility of information (particularly digital information) means that data flows can reveal much of the full expanse of a distributed information ecosystem. These insights can then be leveraged to determine how to prioritize efforts to establish visibility into the organization's cyberinfrastructure. Once sensitive information is identified and the flow of it is mapped through an organization, the other information assets which interact with this information form the base set of assets which are of most significant concern. At this point it will be appropriate to involve relevant stakeholders for these assets and identify sources of information about them.

→ Step 2. Establish visibility.

Asset documentation is only as good as the organization's visibility into the asset being documented. However, maintaining visibility on the full range of information assets can be challenging and often can require large amounts of manual effort to track down information assets. This is compounded by the variety of assets the RCO will be responsible for maintaining. These will typically vary from traditional IT or operational technology assets such as servers, network hardware, HVAC control systems, or desktops, to specialized assets such as supercomputers, telescopes, sensor networks, or specialized instrument control software. To reduce the amount of manual effort required to create and maintain this visibility, organizations should rely on security products and scanners designed to aid with asset discovery and inventory whenever possible. Although these tools typically still require some manual effort (*e.g.*, classifying the assets by appropriate cybersecurity outcomes, *see* **Must 4 (Asset Classification)**), they greatly reduce the total amount of manual effort required.

There are three views, all of which are needed for asset documentation to be successful:

1. Automated from the system itself (agents)
2. Automated from outside the system (scanning)
3. Manual efforts (*e.g.*, asset inventory tags, systems not on the network)

Organizations frequently have preexisting resources that provide effective visibility into their assets and which can be used to populate asset documentation. These include information captured by existing infrastructure such as DHCP logs, configuration management systems such as System Center Configuration Manager (SCCM),⁴⁰ Chef,⁴¹ or Puppet,⁴² capital asset tracking systems, software repositories, and more. These resources should be utilized wherever possible both to

⁴⁰ <https://www.microsoft.com/en-us/system-center>.

⁴¹ <https://www.chef.io/configuration-management/>.

⁴² <https://puppet.com/use-cases/configuration-management/>.

streamline documentation efforts by providing additional information sources and to validate existing information about assets.⁴³

→ Step 3. Develop documentation.

Once organizations have established adequate visibility, it then needs to shape that visibility into a form that is useful for its cybersecurity professionals and decision makers. Since not all information will be relevant, and what information is relevant will vary based on the intended audience, organizations will want to shape what information it documents and how that information is presented.

Organizations should seek examples of common documentation types. There are publicly and commercially available templates and worksheets for common asset documentation types such as network maps, data flow diagrams, and asset inventories.^{44, 45, 46} Trusted CI also has templates available for some common asset documentation types.⁴⁷

Organizations will typically need to aggregate the raw data from data sources identified in step 2 to produce useful documentation. For example, DHCP logs can identify that a device was connected to the network, but it will not provide details such as asset owner or the network services listening on the device. In order to produce a complete picture of each asset multiple information sources will be needed.

Furthermore, organizations should consider documenting assets at various levels of abstraction. It may be practically necessary to characterize types or classes of information (*e.g.*, “personal information on research data participants”), related components (*e.g.*, “custom imaging utilities for processing satellite data”), information systems (*e.g.*, “personnel personal mobile devices”) rather than more discrete units. At the same time, many organizations will have some very specific and even unique assets warranting closer analysis (*e.g.*, a special instrument array; a specific data set). The Trusted CI asset inventory template attempts to strike a balance appropriate for RCOs.

Documentation which describes both assets and how it is protected provides additional utility for operations in a cybersecurity program. A description of how assets are protected provides a valuable tool when determining how a baseline control set applies to specific assets and identifying gaps in existing controls. Where applicable, detailed descriptions of existing security controls are key when

⁴³ Depending on an organization’s chosen baseline control set (*See* **Must 15 (Baseline Control Set)**), these activities may be required or already in place in an organization. Many control sets specify the use of existing information systems to enrich documentation and provide additional context for inventories. For example, in the CIS Controls Version 7.1 subcontrol 1.3 specifies the use of DHCP logs to enrich hardware inventories.

⁴⁴ The Open Science Cyber Risk Profile (OSCRP) provides guidance on assets to consider for science projects. <https://trustedci.github.io/OSCRP/OSCRP.html>.

⁴⁵ The ESNNet Sample Campus & Regional Cyberinfrastructure Plans provide a wealth of relevant examples for RCOs. <https://fasterdata.es.net/campusCIplanning/>.

⁴⁶ Tools such as the Auditscripts CIS Critical Security Controls Version 7.1 Assessment Tools (<https://www.auditscripts.com/download/4229/>) provide a way to both audit and track implementation of security control baselines.

⁴⁷ Trusted CI’s **Information Asset Inventory Template** and **Asset Management Policy Template** may be helpful for RCOs. <https://trustedci.org/framework>

documenting efforts to adhere to regulated data requirements such as Controlled Unclassified Information (CUI) or Protected Health Information (PHI). In addition, this information is often valuable to users of the RCO's CI resources as they often have distinct security or compliance requirements which they must adhere to.

Common Challenges & Recommendations

This section describes some common **Must 3** challenges and offers recommendations on how to overcome them.

→ Network maps, data flow diagrams, and other relevant graphics.

A catalog of information assets based solely on the output of information sources such as network scanners, logs, existing asset inventories, or other sources will provide limited capability to effectively represent important relationships between key components. Organizations should develop a variety of visual aids where practical to better illustrate these relationships. Network maps, data flow diagrams, and other graphics provide an effective and accessible way to capture this information.⁴⁸ These visual aids should capture key elements and interconnections of the systems being described. In some cases it will be necessary for RCOs to create multiple diagrams at different levels of detail according to the intended use case and audience of the visual aids being developed.⁴⁹

→ Level of detail.

The documentation completeness and level of detail must be balanced with the resources available to maintain the entries. An inventory might include a number of details, but at minimum, it should identify the asset, indicate the value or sensitivity of the system via asset category, describe the asset's location, and identify the primary asset owner.

The appropriate balance for an organization will differ by asset type and organization. For example, an organization with a large number of software assets may find that it changes rapidly enough that it is impractical to track and document these changes across the multitude of projects until the organization has developed the capacity to manage the flow of information from repositories and build systems in an automated fashion. In such a case it may be more appropriate to document the class of assets associated with a specific workflow rather than expending manual effort attempting to track changes to the distinct components of the set.

Certain assets are not practical to identify individually. In these cases, it will be more appropriate to identify them collectively by an identified relationship. Virtual assets for example can be ephemeral in nature, which can rapidly lead to documentation of these assets being outdated if captured individually. Documentation which captures the details about the base image and default configurations of these types of assets would more pragmatically serve the same purpose when combined with documentation regarding the hosting environment for the virtual assets.

⁴⁸ https://csrc.nist.gov/glossary/term/network_map.

⁴⁹ For example, the ESNNet Science DMZ Architecture overviews include diagrams at multiple levels of detail to illustrate the principles of the system and distinct use cases. <https://fasterdata.es.net/science-dmz/science-dmz-architecture/>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ Unknown devices.

Any device not in the inventory is “unknown”. Any device on an IP network leaves traces in ARP tables and in network switch interface address tables. Using tools that collect the MAC and IP address information, the data can be correlated with other data sources to sufficiently identify the device and determine if it is authorized and a candidate for inclusion in the inventory or unauthorized and other action may be required. RCOs should look to develop the means to track this information. In cases where an RCO does not have direct access to these resources, the RCO should work to establish the appropriate agreement with responsible parties to establish the appropriate visibility to ensure that unknown devices are identified and accounted for.



Must 4: Asset Classification

Organizations must establish and implement a structure for classifying information assets as they relate to the organization's mission.

Information asset classification is used by an organization to enable the assignment of the organization's information assets into organization-defined categories. The categories include the asset's sensitivity in terms of mission impact and stakeholder requirements. These categories express the types and level of protection required for assets and ultimately are used to aid in control selection and tailoring.

Research cyberinfrastructure operators (RCOs) possess or manage a wide variety of information assets, both for themselves and for projects using its research cyberinfrastructure. Evaluating each asset and developing a unique protection plan is uneconomical. Information asset classification is a tool for grouping information assets with similar desired cybersecurity outcomes⁵⁰ into a manageable number of asset categories which can then be managed as a unit, rather than individually.

Information assets can be characterized in a variety of ways. These include consequences to the mission, information sensitivity, placement within the organization, types of security objectives, and level of abstraction. The RCO then develops categories based on the cybersecurity-relevant impact of compromises of the assets according to the level of concern for the ability to satisfy mission objectives **Must 1 (Mission Focus)** and stakeholder requirements **Must 2 (Stakeholders & Obligations)**.

Why is this a Must?

An RCO needs to understand what types of assets it has as well as assets of research projects using the cyberinfrastructure **Must 3 (Information Assets)** and to document the relative type and importance of those assets. Since RCOs may have to manage specialized assets and complex mission requirements, a classification structure streamlines thinking about the types of assets. The standardization through an information asset classification structure brings significant benefits to the organization, providing a way to efficiently categorize information assets by grouping information assets that have common desired outcomes.⁵¹

The Roadmap

This section describes the steps needed for an RCO to establish and develop its information asset categories. In Step 1, the organization determines the assets from **Must 3 (Information Assets)** relevant to the mission or stakeholder requirements. In Step 2, the organization determines a set of groupings based on assets with similar levels of priority and mission impact. In Step 3, the

⁵⁰ For example, desired outcomes might be resilience, availability, safety, and/or integrity.

⁵¹ See **Must 15 (Baseline Control Set)** and **Must 16 (Additional & Alternate Controls)** for how the categories are helpful in determining the cybersecurity controls to protect the information assets.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

organization defines the categories consistent with the mission to ensure assets are included. In Step 4, the organization refines the asset categories by combining or splitting them as appropriate.

- 🔗 Trusted CI has developed templates that may be appropriate for an RCO to document the results as it is proceeding through these steps.⁵²
- 🔗 The Open Science Cyber Risk Profile (OSCRP) is a joint project of Trusted CI and the Department of Energy's Energy Sciences Network (ESnet).⁵³ The working group developed a "risk profile for open science." The risk profile is a method of categorization of scientific assets and their common risks to science for open science projects.⁵⁴ The core of the document is only about a dozen pages and it was written to expressly avoid technical language and be approachable by PIs and an RCO's management while providing the kind of guidance needed by technical staff to implement cybersecurity controls.

→ Step 1. Determine which assets are of concern for the mission.

The documentation called for by **Must 3 (Information Assets)** is the primary source of input for this step. Presumably, all information assets have some relevance to the mission, but the level of concern about negative consequences serves as a guide to assets to be considered. It is necessary to involve a wide range of stakeholders (*see* **Must 2 (Stakeholders & Obligations)**) to ensure they have been given a voice in the process and critical or important assets are not overlooked.

→ Step 2. Group assets by consequence and type of concern.

The type and severity of potential harm arising from a compromise of the asset to an RCO informs the choice of appropriate grouping. For instance, if there are information assets that process, transmit, store, or control access to data that is not to be viewed by unauthorized parties, it is likely appropriate to be in a "confidential" group. Similarly, there might be a grouping for information assets where it is critical that the data be available. Note that some information assets may have multiple concerns or different priorities with groupings needing to accommodate this fact. For instance a device controller (*e.g.*, for a telescope) might have desired outcomes of integrity, availability, and safety, with availability and safety being considered priority concerns.

→ Step 3. Define categories to cover all assets.

Define categories sufficient to cover all the groupings found in Step 2. The important aspect is to ensure that all the information assets from Step 1 are in a category reflecting the mission impact of the asset and priority. This categorization can be rough since it will be refined in the next step.

⁵² Trusted CI's **Information Classification Policy Template** is available at <https://trustedci.org/framework>.

⁵³ <https://www.es.net/>.

⁵⁴ <https://trustedci.github.io/OSCRP/OSCRP.html>.

→ Step 4. Combine and split asset categories.

Remember that defining categories is intended to **reduce** the effort of managing the information assets. Balance is important. Having too many or too few categories can create an organizational burden that having the categories was supposed to alleviate. In this step, the categories are adjusted to account for real-world situations that improve their useability in subsequent control selection.

Combining categories. Take the case of the previous example of a device controller (*e.g.*, for a telescope) that has desired outcomes of integrity, availability, and safety, with both availability and safety being of critical concern. Another controller might have the same concerns but only safety is a priority concern due to the fact that there are multiple controllers and so availability of a single device is not such an important issue. It may be reasonable to combine these categories since a threat that targeted ALL of the multiple controllers would be critical. In general, assigning an asset to a category with a higher priority concern could result in additional controls being applied, this action must be carefully considered in light of the potentially negative consequences on the mission of the additional controls.

Splitting categories. Although information assets may appear to fall in the same category, there are differences that can call for separate categories: different consequences arising from contractual or legal requirements; different access control mechanisms; and different access control policies. For example, information provided by web servers is likely critical in terms of availability and integrity; however, while some of the data is public, some of the data is restricted to registered users (not a priority concern, but desirable). In such cases, due to the different access policies, use separate categories for the associated assets.

→ Repeat Steps 1 through 4 on a periodic (at least annual) basis.

Due to the dynamic nature of the RCO's environment, the category list should be reviewed and updated to adjust for changes in mission (*see* **Must 1 (Mission Focus)**) or stakeholder requirements (*see* **Must 2 (Stakeholders & Obligations)**), new types of information assets (*see* **Must 3 (Information Assets)**), and changes in the cybersecurity environment. A periodic review (at least annual) allows for a more holistic view of the changes that might be required that were not noticed in more limited incremental updates.

Common Challenges & Recommendations

This section describes some common **Must 4** challenges and offers recommendations on how to overcome them.

→ Explosion in a number of categories.

The goal of creating a classification structure is to create an organizational shorthand for understanding the different profiles of the organization's assets. Determining the appropriate number of asset categories will always be a matter of balancing between total specificity to the needs of a particular asset against the ease of application and generalizability across the organization.

Smaller organizations may lean more heavily on ease of application, employing a small number of categories (as these organizations may lack the resources to manage multiple unique categories). Larger organizations may find greater benefit from employing more categories addressing the particular needs of specific assets.

→ Asset value decoupled from cost.

Distinguishing between mission importance and monetary value when organizations are evaluating whether a particular asset is critical is a common challenge. Although money is often a useful shorthand for importance, an organization's most valuable assets can be relatively cheap or seemingly incidental (*e.g.*, data generated by a sensor array or a safety interlock system to prevent injury to personnel or damage to equipment). A simple strategy to assess criticality is to assume the compromise, failure, or theft of the asset, and evaluate the consequences to the mission or stakeholder requirements. If the mission would fail or be severely compromised, then that asset is critical, and should be categorized in that light.

→ Over-restrictive confidentiality concerns.

While “need to know” is often cited in security literature as a criterion for designing access controls, in a knowledge-based organization such as an open science project it is nearly impossible to predict in advance what data or information is significant to fire the creative processes or critical for some report or proposal. It is more appropriate to use “need to protect” as a guide.⁵⁵ Easy access to data within a research group not only enhances the sense of collaboration but is also a great magnifier of productivity.⁵⁶ (Wikis are a common example of such a collaborative environment.) The requirement to protect the data integrity should not be combined with data confidentiality access controls.

→ Non-mission critical assets become critical as groups.

Assets that are not individually critical can be critical as a group (*e.g.*, a single research workstation vs. all research workstations) relating to the level of abstraction (component vs. system) of the information assets. Non-critical assets are often replicated and therefore relate to an organization's resiliency, the ability of the organization to continue when an asset fails or is compromised. Failures can become critical when the number of failures reaches a certain threshold or when the time to recover exceeds a certain limit.

→ Organizational structure affecting categories.

In more complex RCOs' environments, there may be different governance structures that affect the categorization of assets. Differences in the manner concerns and consequences are viewed, differences in the research projects supported, or differences in stakeholder requirements can all result in distinct categories for parts of the organization. The categories exist to aid in managing the information assets and should be used where they make sense in the organizational structure and to

⁵⁵ Donn B. Parker. 1997. The Strategic Values of Information Security in Business. *Computers & Security* 16, 7 (1997), 576-577. DOI: [https://doi.org/10.1016/S0167-4048\(97\)80793-6](https://doi.org/10.1016/S0167-4048(97)80793-6).

⁵⁶ In one large physics collaboration, the AFS home directories were read-accessible to all authenticated users (excluding email subdirectories).

the people involved. Groups with significant Operational Technology (*e. g.*, an accelerator control group) may be in a different category due to concerns about safety. Assets with similar control requirements but different access control mechanisms and different access policies (*e. g.*, ePHI research data vs. PII) belong in different categories.

→ Missing RCO Categories.

Even though controls may be similar for certain types of business data for the RCO compared with types of regulated research data, it is important to keep categories for research data separate from non-research data since the governance, access control mechanisms and policies are likely to be different. The RCO may not have an appropriate category for research projects where additional protections are required. In that case, the RCO and the research project should jointly decide which has the responsibility for defining the category and implementing appropriate controls.



Governance



Must 5: Leadership

Organizations must involve leadership in cybersecurity decision making.

Organizational leadership includes the senior executives and other decision makers responsible for an organization. These are the people ultimately responsible for the organization who make final decisions regarding the highest priorities. Common leadership roles/titles include Director, Board, Chairman, Chief, Executive, Commander, President, Vice President, Partner, Principal, Owner, Founder, and Secretary. Leaders in these roles are in the best position to adjudicate competing demands for resources across the organization, to include prioritizing cybersecurity.

Like any important programmatic effort that addresses serious risks to an organization's mission, the cybersecurity program benefits from organizational leadership being involved in major decisions and ensuring the cybersecurity program is adequately prioritized (*e.g.*, in terms of resources, budget, and personnel). The impact of cybersecurity risk is no longer confined to the information technology (IT) department and information assets: it's an organization-wide concern.^{57,58} Although organizations naturally include leadership in the event of a major incident, leadership should be involved in other cybersecurity decisions as well, particularly on program resourcing and risk acceptance. This is not to suggest that the senior-most leaders entrench themselves in day-to-day cybersecurity decisions or become technical experts, but rather that they should be involved in the decisions that ensure cybersecurity is in support of the mission.⁵⁹

Because organizational structures and reporting relationships vary, and many research cyberinfrastructure operators (RCOs) sit within parent organizations, RCO leadership may or may not be the people making final decisions for their cybersecurity program. Because the RCO may have distinctive missions, information assets, relationships, and security concerns, RCO leadership may need to fill a security advocacy role in relation to the parent organization in addition to exercising decision making in its areas of discretion and authority (*e.g.*, spending RCO resources on additional controls not provided by a parent institution).

Why is this a Must?

Organizational leaders are the primary agents of the organizations for which they work, representing the organization to the outside world. They are ultimately responsible for an organization, and are best positioned to bear the burdens of tough decisions about risk taking and risk reduction. Organizational leadership is also most frequently and reasonably held accountable for decisions that

⁵⁷ Rothrock, R., Kaplan, J., & Van Der Oord, F. (2018). The board's role in managing cybersecurity risks. MIT Sloan Management Review, 59(2), 12-15.

⁵⁸ <https://hbr.org/2019/11/companies-need-to-rethink-what-cybersecurity-leadership-is>.

⁵⁹ <https://www.weforum.org/agenda/2020/01/global-leaders-must-take-responsibility-for-cybersecurity-here-s-why-and-how/>.

put the organization at serious risk.^{60,61,62} No job roles are more directly and holistically connected with the organization's mission than those of its leadership.

Cybersecurity programs exist to support the organization's mission and must be tailored to achieve this intent, as outlined in **Must 1 (Mission Focus)**. The RCO cybersecurity lead has the responsibility to advise leadership regarding cybersecurity risk (*see* **Must 7 (Cybersecurity Lead)**), and may be responsible for drafting budgets, presenting information about organizational risk, or building business cases for cybersecurity investment. However, RCO leaders ultimately control the allocations of resources, budget, and personnel to support the cybersecurity program.

The Roadmap

This section describes the steps needed for an RCO to involve leadership in cybersecurity decision making. Cybersecurity programs necessarily begin with involvement by leadership with, for example, formative, formal decisions on resourcing and governing cybersecurity activities. This section describes steps RCOs should consider to ensure continued leadership involvement in cybersecurity decision making. These steps follow a typical sequence of: Step 1, assessing the current state; Step 2, identifying gaps and planning for change; Step 3, implementing and experiencing changes; Step 4, evaluating effectiveness and making adjustments as required to maintain an effective level of leadership involvement.

→ Step 1. Assess the current state of leadership involvement.

The first step is to assess the current state of leadership involvement. This should include reviewing organizational policy and associated management processes to assess the roles and responsibilities currently implemented that facilitate cybersecurity decision making, as well as the actual frequency and nature of that involvement. Consider the Musts under the Governance and Resources Pillars as a source to guide the policy and management processes reviews. Specifically, the review should focus on determining roles and responsibilities the organization has in place for cybersecurity risk acceptance **Must 6 (Risk Acceptance)**; resourcing the cybersecurity program to mitigate unacceptable risks **Must 11 (Adequate Resources)**; establishing a budget to maintain the program **Must 12 (Budget)**; and allocating the human capital required for program operations **Must 13 (Personnel)**.

→ Step 2. Assess gaps and plan for change.

Armed with knowledge of the current state from Step 1, this is the phase where the RCO decides the changes it should make to the program, for example, roles, responsibilities, practices, policies, and reporting requirements. The deliverable outcome of Step 2 is an artifact that identifies the current state, the gaps, and plans for how to fill them.

⁶⁰ <https://www.nytimes.com/2017/09/26/business/equifax-ceo.html>.

⁶¹ <https://www.techrepublic.com/article/eight-reasons-more-ceos-will-be-fired-over-cybersecurity-breaches/>.

⁶² <https://www.federaltimes.com/it-networks/2016/02/22/opm-cio-seymour-resigns-days-before-oversight-hearing/>.

Things to consider during this phase include:

- 1) Has leadership been caught by surprise by the results of cybersecurity audits, penetration tests, assessments, or security incidents. Or, either learning about gaps and risks or their actual occurrence too late to be effective in limiting the negative impact?
- 2) How does involvement in cybersecurity decision making compare with involvement in other risk limiting functions of the organization?
- 3) Does leadership know what cybersecurity improvements or changes the cybersecurity lead or other relevant personnel would make if more resources were available?
- 4) If the RCO sits under a parent organization, what role is RCO leadership playing in negotiating and shaping the relationship relative to cybersecurity?
- 5) What are the expectations of funding agencies and regulatory bodies regarding leadership involvement?
- 6) If there is a risk registry for organizational risks, are multiple cybersecurity risks included? (*e.g.*, phishing, ransomware, data corruption)
- 7) Are potential cybersecurity issues considered when signing contracts to acquire funding and other agreements, particularly when access to RCO information resources is involved?

→ Step 3. Implement.

The outcome of Step 3 is the actual experience of implementing new or revised roles, responsibilities, practices, policies, and reporting requirements. If organizational leadership roles and responsibilities have changed substantially, include these in a master information security policy document or similar governing policy.⁶³ If changes are substantial, that is all the more reason to adopt an evaluative perspective (*i.e.*, enter Step 4) early and readily.

→ Step 4. Evaluate and Adjust.

The final step is evaluating the effectiveness of the implementation and making adjustments that help achieve continued leadership involvement. Changes in organizational structure, management processes, or the leaders themselves may require the RCO to revisit one or more of the previous steps. For instance, an organization implementing a new financial management mandate may require the RCO to adjust the processes for cybersecurity budget approval.

Common Challenges & Recommendations

This section describes some common **Must 5** challenges and offers recommendations on how to overcome them.

⁶³ Check out Trusted CI's **Master Information Security Policy & Procedures Template**, available at <https://trustedci.org/framework>.

→ The RCO’s mission results in different cybersecurity needs than those addressed by the parent organization’s cybersecurity program.

Many RCOs are embedded within or part of a larger parent institution, such as a university or shared research facility. The parent organization’s enterprise technical solutions and security policies may not be sufficient to support the diversity and complexity of RCO technology, protect against threats targeted at these research programs, or handle heightened cybersecurity and privacy compliance requirements that may come along with sponsored research.

RCO leadership must consider the burden of providing security beyond the default parent capability. This is the reality when RCOs consider pursuing funding opportunities with more cybersecurity requirements than their parent organization is readily prepared to support (*e.g.*, DFARS Clause).⁶⁴ Together with parent organization leadership, determine the distribution of decision making authority and financial burden.

RCO leadership needs to know where they have discretion to tailor their cybersecurity capability. If this discretion is not sufficiently defined, the RCO should consider researching parent organization policy, seeking guidance from the parent organization’s leadership, and/or consulting with other units with heightened cybersecurity requirements.

→ Who counts as “leadership?”

The definition of leadership presented at the beginning of this chapter lists common titles. It is clear that a company’s CEO, a university’s Board of Trustees, and an RCO’s Director are all “leaders” for their respective organizations. However, many organizations, including many RCOs, have layers upon layers of management, and job titles alone may or may not reflect the actual seniority of the position. For the purposes of cybersecurity decision making and maintaining a competent cybersecurity program, who really counts as leadership? To some extent, this is a decision that needs to be made by the RCO itself. What is essential is the RCO not exclude the most senior, most authoritative roles (*e.g.*, the Director) in cybersecurity decision making.

⁶⁴ DFARS 252.204-7012 <https://www.acq.osd.mil/dpap/dars/dfars/html/current/252204.htm>.



Must 6: Risk Acceptance

Organizations must formalize roles and responsibilities for cybersecurity risk acceptance.

Risks are uncertain events or conditions—such as a successful cyber attack—that, if they occur, have a positive or negative effect on the organization’s mission.⁶⁵ Risk acceptance is a decision to acknowledge a risk and not take further action unless the risk occurs. Organizations apply a variety of strategies to manage risk, but decisions to accept risks are of central importance and complexity in cybersecurity. Formalization of roles and responsibilities means documenting them in organizational policy and using them to guide delegation of authority and accountability.

Cybersecurity is a complex decision making discipline, and no decision is more important than when **not** to invest in more cybersecurity.

Organizations have to deal with all kinds of risks, and have a variety of strategies available to them to do so. Various authorities define these strategies in slightly different ways, but they typically include things like risk avoidance, risk acceptance, risk transference,⁶⁶ risk escalation,⁶⁷ and risk mitigation (aka treatment).⁶⁸ At the core of any risk management approach to cybersecurity governance, however formal or informal, is **risk acceptance**.⁶⁹ Risk acceptance is a decision to acknowledge a risk and not take further action unless the risk occurs. Risk acceptance decisions can be temporary (in anticipation of future mitigation) or long-lived. One flavor of risk acceptance involves acknowledging a risk and doing nothing to avoid, transfer, or mitigate it. Another form of risk acceptance is to undertake some risk mitigation, but decline to go to extremes in mitigating cybersecurity risk. In these cases, some degree of *residual risk*⁷⁰ acceptance accompanies risk mitigation. For example, to protect myself in the event of a motorcycle accident, I am willing to pay for and put up with the discomfort of wearing a good quality helmet. However, I’ve declined to purchase and wear motorcyclist body armor. As such, I’ve accepted the (quite substantial) residual risk of serious bodily injury or death against which my helmet cannot protect.

This chapter focuses on the distribution of this type of cybersecurity responsibility in research

⁶⁵ Derived in part from Project Management Institute, Inc., (2017). A Guide to the Project Management Body of Knowledge (PMBOK® Guide). 6th ed. Newtown Square, PA. *See also*, National Science Foundation, Major Facilities Guide, NSF 19-68, September 2019.

⁶⁶ A common form of risk transfer is buying insurance, *e.g.*, buying car insurance to cover the risk of vehicular damage. Cyber insurance is a complex topic, and the market is still evolving. While not irrelevant to RCOs, we note that many cyber risks to RCO missions cannot be managed through transfer. For example, no amount of insurance can bring back lost opportunities to observe fleeting or time-sensitive phenomena.

⁶⁷ Particularly relevant to RCOs with parent organizations, risk escalation is discussed at the end of **Must 14 (External Resources)**.

⁶⁸ Project Management Institute. (2017). A Guide to the Project Management Body of Knowledge (PMBOK® Guide). 6th ed. Newtown Square, PA; Schwalbe, K. (2019). Information Technology Project Management. 9th ed. For more discussion on the language of risk, *see* the Risk Terminology discussion later in this chapter.

⁶⁹ A complete treatment of risk management approaches is outside the scope of this Framework Implementation Guide. We discuss the value and complexities of risk management approaches to cybersecurity later in this chapter.

⁷⁰ *I.e.*, the risk remaining after controls are implemented.

cyberinfrastructure operators (RCOs). Although risk acceptance can happen informally at all levels of an organization (*e.g.*, any time personnel click links in emails that **could** be phishing attacks), formal risk acceptance is a structured process whereby the organization assigns individuals with authority and/or accountability for certain risks, and those individuals document and communicate their decision making on those risks. This formalization is particularly important for decisions with the potential for high impact on mission, resources, and basic functioning of the organization. These formal risk acceptance responsibilities naturally fall upon organizational leadership, and those leaders may choose to retain or delegate risk acceptance responsibilities to other actors in their organization.

Trusted CI recommends that leadership (*e.g.*, the RCO director, university senior executives) retain risk acceptance responsibility for major governance and resourcing decisions.⁷¹

Why is this a Must?

- 1) Without formalization of risk acceptance roles and responsibilities, particularly regarding what decisions to retain at the highest echelons of an organization, a great deal of cybersecurity risk acceptance may be left to people not in the best position to make the judgment calls. For RCO's these include diverse researchers and collaborating organizations that interface with the RCO's information assets in a great variety of ways.
- 2) Informed risk acceptance supports risk taking, and risk taking is what makes great things happen. Great scientific discoveries happen because someone took a risk, with money, with life and limb. However, uncontrolled or *ad hoc* risk taking can put the entire organization in jeopardy, where unreasonable risks are accepted because the actors in the best position to make those decisions were unaware or underinformed. Formalizing these roles and responsibilities helps to get people asking the right questions and to ensure that the right people are making the important decisions.
- 3) Human judgment and communication play a big role in deciding how much mitigation is enough mitigation. Contemporary organizations need a particularly thoughtful approach to deciding who will make the call that there is enough cybersecurity. Well-informed risk acceptance is the ideal, but due diligence is challenging for several reasons:
 - a) After reasonable mitigations are applied, there may be a great deal of residual risk, and gauging how much is extremely challenging. There are many different types of inputs to cybersecurity decision making, each with potential utility and costs.⁷²
 - b) Much of cybersecurity risk stems from intelligent adversaries working against difficult-to-secure hardware, software, and human activities.
 - c) Cybersecurity's scariest risks are "black swans,"⁷³ *i.e.*, high impact, but very infrequent events for any given organization, and difficult to predict in terms of timing. "Risk acceptance is the least expensive option in the near term and often the most

⁷¹ See below, Roadmap, Step 2.

⁷² See later discussion on Inputs to Risk Acceptance Decision Making.

⁷³ See the classic text, Taleb, *The Black Swan: The Impact of the Highly Improbable*, Random House, New York, 2010.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

expensive option in the long term should an event occur.”⁷⁴ This is especially true for black swans.

- d) There appears to be no end to theoretically possible steps an organization could take to reduce cybersecurity risk (given unlimited resources and unlimited tolerance for inconvenience). Decision makers face a persistent question of “how much security is enough?”

The Roadmap

This section describes the steps needed for an RCO to formalize roles and responsibilities for risk acceptance. Trusted CI recommends that the RCO determine the following: Step 1 is to determine the current state of formalization of risk acceptance roles and responsibilities; Step 2, is to determine what responsibility to formally retain with leadership and what to delegate; Step 3 is to embody these decisions in policy, job descriptions, and other artifacts and processes; and Step 4 is to monitor, evaluate, and be willing to revisit decisions.

→ Step 1. Determine the current state.

Whether the RCO is new, newly addressing risk acceptance roles and responsibilities, or readdressing these in light of experience and this guidance, addressing this Must requires some discovery. What is the current state? In the absence of prior formalization, the distribution of risk acceptance decisions may be not only highly informal, but distributed widely. Even with prior formalization, it is possible that policy and practice have diverged. Important questions to ask in this phase are:

- a) What, if any, roles are leaders playing in risk acceptance decisions?
- b) What are the formal venues, if any, where key players discuss cybersecurity risks and make decisions?
- c) What risks to the RCO’s mission are, in effect, being accepted by parties who do not work directly for the RCO or personnel in very junior or temporary roles?

If the RCO is already working with a set of formalized risk acceptance roles and responsibilities, also consider the questions below in Step 4.

→ Step 2. Decide which risk acceptance responsibilities to retain with senior leadership, and which to delegate.

The importance of leadership involvement in cybersecurity decision making is central to the Framework, and is the subject of **Must 5 (Leadership)**. This chapter follows immediately for a reason: decisions regarding which risk acceptance responsibilities to retain with leadership structure the collaboration, communication, and judgment calls that form the foundation of a functional, competent cybersecurity program. Leaders have to lead organizations, but they also have to delegate to be effective. Here are a few basic guiding principles to aid these decisions:

⁷⁴ Risk Mitigation Strategy Development Susan Snedaker, Chris Rima, in Business Continuity and Disaster Recovery Planning for IT Professionals (Second Edition), 2014. <https://www.sciencedirect.com/topics/computer-science/risk-acceptance>.

- a) Retain risk acceptance responsibility for major governance and resourcing decisions with RCO leadership. These decisions would include approving major budget increases or decreases, hiring the cybersecurity lead, and adopting the core policies of the cybersecurity program.
- b) Retain risk acceptance responsibility for a decision not to address a Must or to address a Must with minimal resources with RCO leadership. For instance, only senior leadership are appropriately placed to forgo a review of legal requirements in addressing **Must 2**, to allocate no personnel resources to cybersecurity under **Must 13**, or to forgo the adoption of a baseline control set under **Must 15**.
- c) Delegate risk acceptance responsibility to positions that are (a) incentivized to ensure the continuity and success of missions and business functions supported by the relevant information assets and (b) have sufficient cognizance and control over the information assets to be a meaningful part of cybersecurity risk communications and action-taking. Some organizations use the concept of an “asset owner” or “steward” where particular roles are delegated general responsibility for an information system, type of information, or service.
- d) Take care when deciding to delegate risk acceptance responsibilities to roles that are heavily focused on risk-taking or risk-reduction. Roles sitting at the further extremes are a natural part of organizations of any size or complexity, but they require greater oversight to keep both opportunities and risk reduction in a balanced perspective. We discuss CISOs and risk acceptance later in this chapter.

→ Step 3. Embody these decisions in policy, job descriptions, meeting agendas, and other relevant artifacts and processes.

Risk acceptance roles and responsibilities can be formalized and communicated in a variety of ways: in policy, in job descriptions, in meeting minutes, and any number of other artifacts. In formalizing these roles and responsibilities, RCOs should include not only who is responsible for making risk decisions for a particular asset, asset type, workflow, or environment, but also the associated communication and reporting responsibilities. Organizations should include reporting responsibility in proportion to the degree of decision making delegation, and risk acceptors should rely heavily on cybersecurity and other risk-reduction roles (*e.g.*, general counsel) to advise on the risk environment. For major risk acceptance decisions, RCOs should go beyond formalizing the roles and responsibilities, and formally document the decisions themselves (*e.g.*, in risk registers).⁷⁵ As a general heuristic, the level of formality required for a particular risk acceptance decision should correspond with the magnitude of organizational risk, with the most impactful and/or likely risks requiring the most procedure and/or documentation.

→ Step 4. Monitor, communicate, and be willing to adapt.

As with all the Framework’s Musts, an important step involves evaluating whether things should change. Governance decisions can seem hard to change. Revising these decisions requires leadership time and involves shifting organizational power and burden. Because cybersecurity is a relatively

⁷⁵ See, *e.g.*, National Science Foundation, Major Facilities Guide, NSF 19-68, September 2019, Risk Identification and Risk Register (6.2.6.2).

young and dynamic discipline, it is particularly important that organizations remain open to changes in cybersecurity roles and responsibilities, including the delegation and distribution of risk acceptance. At predictable intervals (*e.g.*, annually or every other year), RCOs should evaluate the following questions:

- a) Have those roles charged with cybersecurity risk acceptance been actively involved in communication with cybersecurity personnel and making decisions? If not, are there institutional or other barriers standing in the way?
- b) Do risk acceptors feel adequately informed as to opportunities to reduce cybersecurity risk?
- c) Are risk acceptance decisions sufficiently well-documented that leadership can review them?
- d) Are any risk acceptors having their decisions overridden on a regular basis? If so, is there a gap between policy and practice? These gaps can be closed via policy enforcement or policy revision, but they should be closed.
- e) If leadership has retained a great deal of responsibility in this area, have past leadership decisions substantially set precedence and shaped organizational culture such that more can be delegated?

Common Challenges & Recommendations

This section describes some common **Must 6** challenges and offers recommendations on how to overcome them.

→ Inputs to risk acceptance decision making.

There are many different types of informational inputs that can be used to inform cybersecurity decisions about investing in more security or saying “that’s enough for now.” They can inform big programmatic decisions, point-in-time judgment calls, or both. They can be tailored specifically to your organization or community, or be very general. They can address particular threats or types of threats, or (again) be very general. They can be focused on a basic standard of care (*e.g.*, what is good enough) or how to be extra safe.

This area presents a challenge for at least three reasons: (a) Each of these input types come with costs (*e.g.*, money, time, effort to understand and translate to decision makers); (b) they each have both utility and limitations; and (c) as put into practice, each of these input types can be done well or poorly.

These input types include, but are not limited to:

1. lessons learned from cybersecurity incidents;
2. logs and outputs of monitoring systems;
3. threat reports and bulletins from third parties;
4. results of various types of cybersecurity assessments, including, gap analyses against a baseline control set (*see Must 15*);
5. the results of risk assessments/analyses;
6. analysis from third party or parent organization security operations centers (SOC);
7. benchmarking information and reports;

8. collections of best practices (including those found in compliance requirements); and
9. expert briefings.

Because no risk acceptor or cybersecurity expert can gather and fully weigh the value of all the possible inputs, here are a few general guidelines for what inputs to prioritize:

Prioritize inputs that:

1. have particular relevance to your own local environment and community;
2. provide information about what is already happening or imminently about to happen, versus what might possibly happen in the future;
3. come from trusted sources;
4. provide a clear evidentiary basis for their assertions and estimations;
5. provide actionable information and guidance; and/or
6. are independently corroborated by other informational sources.

The “and/or” and item 6 are important. No single source of input into risk decisions is the one to rule them all.

→ What about the CISO accepting risk?

Having the RCO’s cybersecurity lead accept cybersecurity risk should only be done after careful consideration. In a world of limited resources and hard-to-understand risks, CISOs should and frequently have a need to advocate for cybersecurity and provide a perspective on worst case scenarios and what more could be done to reduce cyber risk to mission. Even extremely business-minded and mission-savvy CISOs fulfill a primary function much like that of a lawyer to a client or doctor to a patient. Generally, cybersecurity personnel are positioned to provide advice and services rather than to make final, high impact decisions on how much security is too much security.

Areas that make the most sense for CISOs to accept cybersecurity risk include day-to-day tactical decisions within the scope of the cybersecurity services that the CISO/cybersecurity team are providing. Organizational leaders may need to approve a set of controls and strategies to be used; within the implementation of those controls and strategies CISOs and their personnel need to make point-in-time decisions using their knowledge, research, and judgment.

→ The complexity of risk management approaches to cybersecurity.

Formal risk management approaches can be valuable, but also bring substantial procedural and communications learning requirements to an organization. National Institute of Standards and Technology (NIST) Special Publication 800-39 (Managing Information Security Risk: Organization, Mission, and Information System View)⁷⁶ captures the potential value and complexity of risk management:

⁷⁶ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-39.pdf>, pp. 1-2.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

“Managing information security⁷⁷ risk, like risk management in general, is not an exact science. It brings together the best collective judgments of individuals and groups within organizations responsible for strategic planning, oversight, management, and day-to-day operation—providing both the necessary and sufficient risk response measures to adequately protect the missions and business functions of those organizations.”

“The complex relationships among missions, mission/business processes, and the information systems⁷⁸ supporting those missions/processes require an integrated, organization-wide view for managing risk.”

Capturing this complexity in formal language and processes requires effort and education. Part of the complexity, and resulting learning curve, is linguistic. There are a number of different publicly available risk and risk management lexicons (some attached to risk management frameworks and some not), and they do not offer uniform definitions of the core concepts discussed in this chapter (including “risk acceptance”). Either through experience with these or in informal usage, people may use the same terms inconsistently or different terms to mean the same thing.⁷⁹

The Trusted CI Framework’s underlying philosophy is based on AFCEA’s Economics of Cybersecurity.⁸⁰ Rather than promoting the use of a fully robust and complex risk management approach (*e.g.*, ISO 31000 or NIST SP 800-39⁸¹) this Framework relies on the AFCEA’s proven assertion that a great deal of cybersecurity risk can be avoided or mitigated by adopting and thoughtfully using a quality baseline control set (*see Must 15*). This greatly reduces the number of judgment calls that remain.

Before investing a great deal in full adoption and implementation of a formal risk management methodology, we recommend the following:

1. Do not adopt more risk terms than you really need to have productive conversations with actionable outcomes.
2. Whatever terms you do use, select an explicit operational definition and make sure relevant people are aware of the terms and what they mean.
3. Consider whether formalizing risk acceptance roles and responsibilities, along with addressing the other Musts in the framework, is sufficient to reduce cybersecurity risk to acceptable or optimal levels.

⁷⁷ Reminder: The Trusted CI Framework does not draw a strong distinction between “information security” and “cybersecurity.”

⁷⁸ Reminder: Information systems are included in the Trusted CI Framework’s Definition of “information assets”.

⁷⁹ In our research for this chapter, we looked at the definitions found in common cybersecurity standards (ISO, NIST), in the academic risk management literature, and in project management authorities (PMI, Schwalbe). We did not find a third party lexicon that we felt was straightforward enough to adopt as the default for this Framework, or one that we felt confident we should recommend to RCOs as “the best.”

⁸⁰ <https://www.afcea.org/committees/cyber/documents/cybereconfinal.pdf>.

⁸¹ <https://www.iso.org/iso-31000-risk-management.html>.



Must 7: Cybersecurity Lead

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. This position, often referred to as the Chief Information Security Officer (CISO), ensures the program educates and advises decision makers on cybersecurity matters, including risk identification and mitigation, and policy development. The position also provides leadership for services like incident response coordination, and cybersecurity control selection and monitoring.

Research cyberinfrastructure operators (RCOs) must establish a role with the primary responsibility for the cybersecurity program and its day-in-day-out management, including communicating cybersecurity risks and potential mitigations, directing cybersecurity staff, and overseeing control implementation and other cybersecurity operations. The role should have authority to advise and plan and—after management approval—implement the needed cybersecurity controls, mitigations, policies, and practices. In some RCOs, this individual may be a primary liaison to the parent organization’s cybersecurity team. The role may justify a full-time position depending on the mission, size, and complexity of the RCO; many details depend on the organization’s needs.

Trusted CI recommends establishing a cybersecurity lead role that reports directly to RCO leadership.⁸²

Ideally, a cybersecurity lead is knowledgeable in cybersecurity generally and understands specifically how information assets relate to the organization’s mission. These cybersecurity leaders can effectively communicate the issues and trade-offs of different strategies for providing adequate information asset protection. Their skills facilitate informed decisions by RCO leadership **Must 5 (Leadership)** and risk acceptance decisions **Must 6 (Risk Acceptance)**.^{83,84,85}

In some cases, RCO or parent organization leadership may retain RCO cybersecurity functions with the parent organization or outsource to a third party. Even in these cases, we recommend that the RCO explicitly identify a cybersecurity lead role to advise RCO leadership, manage any RCO-specific cybersecurity services, and assist in a liaison role with the parent organization and any third party providers.

⁸² Discussed in The Roadmap Step 1 and Common Challenges & Recommendations: Cybersecurity Lead Reporting to the CIO.

⁸³ <https://itchronicles.com/security/5-qualities-of-a-great-ciso/>.

⁸⁴ https://www2.deloitte.com/content/dam/insights/us/articles/ciso-next-generation-strategic-security-organization/DR19_TheNewCISO.pdf.

⁸⁵ <https://securityintelligence.com/the-expanding-role-of-the-ciso-seven-attributes-of-a-successful-security-leader/>.

Why is this a Must?

Given the complexity and broad organizational impact of cybersecurity issues, coordinated decision making and execution requires a dedicated role to address those issues. Moreover, having a cybersecurity lead reflects the importance the RCO places on cybersecurity and facilitates communication concerning cybersecurity obligations with internal and external stakeholders.

The Roadmap

This section describes the steps needed for an RCO to establish a lead role with responsibility to advise and provide services to the organization on cybersecurity matters. These roadmap steps are more straightforward for new RCOs or established RCOs just beginning formal program development. Step 1 is to determine the reporting line of the cybersecurity lead. Step 2 is to develop a job description. Step 3 is to hire or promote a person. Step 4 is to evaluate the effectiveness of the cybersecurity lead role and adjust the role as necessary.

→ Step 1. Determine the Reporting Line for the cybersecurity lead role.

The reporting structure affects the kinds of skills and knowledge of the best person to fill the cybersecurity lead role, (*e.g.*, manager, technology expert, communicator), and so influences the job description in Step 2. As cybersecurity increasingly is a critical organization-wide function, senior leaders need advice on the cybersecurity implications of their decisions. Many organizations are moving toward having the cybersecurity lead report directly to the organization's head or at least to a leadership member.⁸⁶ Reporting directly to a senior manager of operational activities (Chief Operations Officer or equivalent) may provide needed operational support. The RCO might also have the cybersecurity lead report to an organizational unit that has a risk mitigation function such as legal counsel, a financial officer (CFO), or Chief Risk Officer (CRO).⁸⁷

→ Step 2. Write a Job Description.

There is no shortage of example cybersecurity lead job descriptions available online. Some are relatively brief, and others are pages long. However, in reviewing the similarities across many, the five critical elements of the cybersecurity lead job description are straightforward:⁸⁸

- Develop and maintain an RCO-wide cybersecurity program.
- Communicate regularly and effectively to management and stakeholders to help leadership understand risks versus benefits of decisions that impact cybersecurity.
- Develop and implement a plan to attract (internally or externally), train, and retain individuals with cybersecurity expertise, capability or interest in learning.
- Participate in information sharing organizations and react to threats.
- Develop and maintain an incident management structure.

⁸⁶ <https://www.csoonline.com/article/3278020/does-it-matter-who-the-ciso-reports-to.html>.

⁸⁷ *Ibid.* Contains a list of various CISO reporting possibilities and lists the advantages of each option.

⁸⁸ <https://securityintelligence.com/the-five-most-critical-tasks-in-the-ciso-job-description/>.

A comprehensive list of knowledge, skills, and abilities to potentially include in a job description for a cybersecurity lead position is available from US-CERT.⁸⁹ A recruiting firm has provided a template for a CISO job description that aligns well with our guidance.⁹⁰ In an RCO, a cybersecurity lead will need to have an understanding of a research and education organization.

→ Step 3. Hire or Internally Appoint a Cybersecurity Lead.

The selection of a cybersecurity lead is very much dependent on organizational factors and the pool of candidates. Due to the shortage of experienced professionals in the field, it may be necessary to perform several searches or reset expectations before finding a suitable candidate. As another option, selecting the cybersecurity lead from inside the RCO can be useful because the candidates are already familiar with the RCO's missions, culture, and environment and can "grow into" filling the requirements of the position with experience and professional development.

→ Step 4. Evaluate Performance and Adjust Role as Necessary.

Changes in organizational size, complexity, obligations, or other mission elements may affect performance and result in the desire for changes to the role or reporting structure. As with all programmatic decisions (*see* **Must 10 (Evaluation & Refinement)**), the scope of duties and reporting structure of the cybersecurity lead should be evaluated periodically to address shifts in organization composition.

Common Challenges & Recommendations

This section describes some common **Must 7** challenges and offers recommendations on how to overcome them.

→ Cybersecurity Lead Reporting to the CIO.

It is not uncommon for RCOs (and contemporary organizations generally) to have a cybersecurity lead report to a Chief Information Officer (CIO). New RCOs should be aware of the potential problems with this reporting arrangement. RCOs already using this reporting structure should consider that added complexities in the environment, new threats, new stakeholder obligations, and new people may warrant a change in the reporting structure.

While there is no hard-and-fast rule to whom the cybersecurity lead should report, there are downsides to reporting to the CIO. The CIO may have different incentives than the cybersecurity lead, and the result can be a conflict for resources. Brian Contos, CISO for Verodin, has said, "Managing security effectiveness and risk management transcends IT and has to operate at an executive level so that technical and non-technical decision makers can be armed with evidence based data in order to make business decisions more effectively and efficiently from an informed position."⁹¹

⁸⁹ <https://niccs.us-cert.gov/workforce-development/cyber-security-workforce-framework/executive-cyber-leadership>.

⁹⁰ https://images.idgesg.net/assets/2018/01/the_ultimate_ciso_job_description_ebook.pdf.

⁹¹ <https://www.csoonline.com/article/3278020/does-it-matter-who-the-ciso-reports-to.html>.

Several articles cite reasons to avoid CISO-to-CIO reporting, including:

1. “Security is an issue for the entire company, not just the IT department. As a CISO ‘A CISO's job is not to protect IT’ - a CISO's job is to protect the business.”⁹²
2. “Organizations where CISOs report to CIOs have 14% more downtime due to security incidents, according to a study by PwC.”⁹³
3. “Organizations where the CISO reports to the CIO have financial losses that are 46% higher, according to the same PwC research.”⁹⁴

Those downsides hold even for organizations that have elevated the CIO to RCO leadership teams and are likely only worsened where the CIO is multiple steps below that level.

The RCO's mission and relevant technologies (including operational technology, *e.g.*, sensors, accelerators, telescopes) may fall outside of the IT umbrella. Research projects may have requirements to develop or install a specialized computing infrastructure. A cybersecurity lead and cybersecurity program separate from the RCO's IT organization may better position them to address the full scope of the organization's mission. The cybersecurity lead must still partner with other units in support of the mission (*see* **Must 1 (Mission Focus)**) and not allow security to become an end in itself.

→ Combining Physical Security and Cybersecurity Roles.

RCOs may want to consider aligning the cybersecurity function with other security functions (*e.g.*, physical, safety) or risk-reducing functions (*e.g.*, legal). Aligning cybersecurity with the security chain of command under a Chief Security Officer (CSO) can help avoid siloing of related security expertise. Close organization grouping of security disciplines can mitigate the risk of intra-organization competition for limited security resources. In the modern world, cybersecurity overlaps with physical and personnel security and public safety. In many cases, physical and personnel security have a common reporting structure. The difference with cyber has, in the past, been the heavy technology focus. As technology becomes a more significant part of physical and personnel security along with privacy issues there is a stronger argument for a closer relationship between all the security functions. Cyberattacks and technologies increasingly transcend the cyber and physical domains, both in their execution and their effects.

→ Organizational Isolation of Cybersecurity.

Issues with organizational isolation of cybersecurity responsibilities can be twofold: (1) the cybersecurity team can view security solutions without regard to how the security is affecting the mission; (2) other parts of the organization have the view that cybersecurity is someone else's problem and not engage in localized security. With mature domains such as Human Resources and Finance, an RCO with a number of organisational units will have personnel in each unit with Human Resources and finance responsibility. These personnel do not report directly outside the unit but still see that the unit follows the policies and procedures of their domain. Similarly, a unit cybersecurity

⁹² <https://www.linkedin.com/pulse/why-ciso-should-report-directly-ceo-cio-prince-rana/>.

⁹³ <https://www.linkedin.com/pulse/cisos-should-report-cios-don-welch>.

⁹⁴ *Ibid.*

coordinator would have responsibility for awareness and program implementation at the unit level. The coordinator would also provide feedback when cybersecurity policies are negatively impacting the unit's ability to perform its mission. Additionally, including cybersecurity acumen as part of individual performance appraisals helps to increase awareness that the responsibility for cybersecurity extends to everyone.



Must 8: Comprehensive Application

Organizations must ensure the cybersecurity program extends to all entities with access to or authority over information assets.

The entities may be either individuals or organizations. Access includes the logical or physical ability to view, create, modify, or destroy information, or modify or destroy information systems. Authority includes legal, administrative, logical, or physical control of information assets.

The Comprehensivity Principle⁹⁵ states “identify and account for **all** relevant systems, actors, and risks in the environment.” This Principle is highly relevant to **Must 8**, as it similarly requires organizations to identify and account for all entities that should be covered by the cybersecurity program. While no cybersecurity program can manage *all* sources of risk, it can and should extend to all entities with access to or authority over the organization’s information assets.

Organizations can use various means to extend the cybersecurity program’s coverage, including technical controls, contractual clauses, employee and contractor agreements, and appropriate-use or other terms-of-service agreements. Policies, procedures, and controls facilitate the execution of the cybersecurity provisions in those contracts or agreements.

Why is this a Must?

If the cybersecurity program fails to reach and appropriately cover all entities with access to and authority over information assets, the organization is vulnerable to compromise. Organizations increasingly have highly distributed user bases, customers, and stakeholders. Research cyberinfrastructure operators (RCOs) provide services to third parties, rely on third party services and technology, and collaborate in myriad ways to accomplish their missions. In addition to accounting for traditional organizational personnel and units, RCOs need to consider these more complex third parties. The cybersecurity program should cover these services.

Whether malicious, accidental, or environmental in origin, gaps or lapses in program coverage can result in a variety of cybersecurity issues. For example, visiting scientists sometimes bring computers from their home institution and place them on the RCO’s network. If the home institution no longer patches the computers, the unpatched vulnerabilities result in a risk for the RCO. Moreover, some compliance regimes (*e.g.*, HIPAA security and privacy rules) require everyone to have access to and understand organizational policies and procedures. Implementing these requirements is possible only if the program includes all relevant entities.

⁹⁵ See <https://cacr.iu.edu/about/principles.html>.

The Roadmap

This section describes the steps needed for an RCO to ensure the cybersecurity program extends to all entities with access to and authority over the organization's information assets. This process involves many other Musts in this Framework.⁹⁶ Implementation progress with **Must 2 (Stakeholders & Obligations)**, **Must 3 (Information Assets)**, **Must 4 (Asset Classification)**, and **Must 15 (Baseline Control Set)** contribute significantly to the ease of implementing this Must. Step 1 is to identify the entities and their access to or authority over information assets. Step 2 is to document the gaps in existing program coverage: in contracts and agreements; in controls, and in policies and procedures. Step 3 is to develop prioritized plans for extending protective measures. Step 4 is to implement the plan as resources allow, evaluate its success, and modify the plan as necessary.

→ Step 1. Identify and Classify Entity Access and Authority Roles.

There are a number of sources that RCOs can analyze to help identify entities with access to and authority over information assets. Particularly important sources to analyze are current contracts, service level agreements (SLAs), grant material, employment contracts, internal policies, rules and regulations, and other agreements as they apply to the organization's information assets. Using these sources, RCOs can compile a list of entities that the cybersecurity program should extend to.

Depending on the number of entities identified, the RCO may want to categorize internal and external entity roles. A solid breakdown of roles can include administration, technical support, cybersecurity personnel, staff researchers, contractors, vendor support personnel, visiting researchers, and the general public. An RCO may want to select a subset of those roles as a starting point when evaluating access and authority over information assets.

Finally, RCOs may also find it useful to subdivide these roles by defining more specific access or authority categories. For instance, RCOs can subdivide the "vendor support personnel" category by the specific assets they support.

→ Step 2. Document Gaps in Existing Program Coverage.

A good starting point for finding and documenting gaps is to review 1) contracts and agreements with third parties and 2) internal controls, policies, and procedures.

⁹⁶ An RCO's cybersecurity program touches **Must 2 (Stakeholders & Obligations)**, **Must 15 (Baseline Control Set)** and **Must 16 (Additional & Alternate Controls)** regulate the ability of information system users to perform actions on **Must 3 (Information Assets)** beyond their access rights or authority consistent with the **Must 1 (Mission Focus)**. These controls supplement **Must 9 (Policy)** that define the access to or authority over the information assets, how and to whom the different controls apply, and establish the conditions for granting and revoking access to or authority over those assets. **Must 7 (Cybersecurity Lead)** provides advice on these entities' effect on the cybersecurity program. **Must 5 (Leadership)** has a role in considering the cybersecurity program's impact when contemplating organizational changes or entering into contracts or agreements. Leadership also provides resources through **Must 12 (Budget)**, **Must 13 (Personnel)**, and **Must 14 (External Resources)**, ensures **Must 11 (Adequate Resources)**, and **Must 6 (Risk Acceptance)**.

Document gaps in contracts and other agreements.

Identify where explicit or implied use of the organization's information assets from Step 1 does not include language extending the cybersecurity program's coverage. By working with RCO or parent organization resources (*e.g.*, legal counsel, human resources, procurement, contract administration, and research project management), discover and document where clauses that define roles and responsibilities are missing from contracts or agreements with partnered entities.

Document gaps in controls, policies, and procedures.

Identify the extent to which existing controls, policies, and procedures from Step 1 apply to the various entities' roles, access rights, and authorities; determine and document the gaps. This analysis should bring to light missing or inadequate controls and identify security protections mismatched to an asset's function and purpose.

→ Step 3. Develop a Strategy for Ensuring Program Coverage.

Next, the RCO should develop a strategy to address gaps within the program's implementation, identified in Step 2. By working with organizational units that monitor or enforce program compliance (*e.g.*, contracts and grants department, legal counsel, human resources, internal audit), RCOs can negotiate to include appropriate language in contracts and other agreements. Establishing or fostering relationships with internal units can develop a strategy to implement controls, policies, and procedures extending to the RCO's information assets.

A concrete plan derived from the above strategies can now be developed. The plan can help convince decision makers to provide resources (*see* **Must 11 (Adequate Resources)**) to address gaps that present a risk to the RCO.

→ Step 4. Implement the Plan and Evaluate Progress.

Finally, the RCO should assign responsibility for implementing and monitoring the progress of the plan developed in Step 3. Examples of implementation and monitoring activities include:

- Integrate technical control solutions like Identity Access Management (IAM) and Federated Identity Management (Federated IdM) to facilitate external user access in a controlled manner.
- Implement monitoring solutions like centralized logging of information asset access and other system events.
- Develop onboarding and offboarding policies and procedures for internal and external entities. Ensure the **Must 9 (Policy)** lifecycle is followed.
- Implement monitoring capabilities and technical restrictions for vendor support personnel.

Evaluate progress while implementing the plan and make adjustments as necessary.

Common Challenges & Recommendations

This section describes some common **Must 8** challenges and offers recommendations on how to overcome them.

→ Mismatch Between RCO Policies and those of Research Groups or Other Entities.

There can be a mismatch or conflict between RCO policies and the policies of research groups or other entities (*e.g.*, vendors or federated organizations)⁹⁷ with access to or authority over organizational information assets. RCOs should determine if a separate policy section or subsection is necessary. An additional policy may be required to address unique situations so that research can proceed without undue risk. It may be necessary to create a policy that outlines the minimum requirements for external entities to participate in the organization's infrastructure. Such a policy would constrain the entities to the assets they need, like a dedicated IP address space, VPN groups, local affiliate accounts, or accounts integrated with a federated IdM. Include legal counsel, human resources, procurement, contract administration, and research project management when developing policy regarding external entity responsibilities. For instance, a research group might bring its own computing infrastructure and systems with special scheduling requirements for scanning or patching.

→ Increased Complexity Due to Internal and External Influences.

An RCO may be embedded within a larger parent organization or have relationships with external entities with access to or authority over RCO assets. In those cases, the RCO may contend with increased diversity and complexity in technical controls for authority (*e.g.*, assigning appropriate roles and permissions). The RCO must evaluate how to tailor controls to ease administration, minimize the complexity of changes, and automate maintenance. For instance, an RCO might normally require a centralized account for access to information assets, but a research group might need to provide access to information to collaborators without requiring the overhead of enrolling in a centralized account scheme.

The RCO should plan for and implement an access management approach that balances stakeholders' access and authority needs with the fewest operational barriers for the RCO to fulfill its mission. For instance:

- Consider the underlying infrastructure that supports access management; depending upon the RCO's placement in the organization, central access control models vs. local/distributed access control can either streamline or complicate access control.
- Balance the protection of the RCO's information assets with enough flexibility or modularity to support business continuity and research activities.

⁹⁷ Note that WISE has developed simple common policies for federated research infrastructures which comply with some general agreed trust framework (like SCI). <https://wise-community.org/sci/> The aim is to minimize conflicts with local site and other stakeholder policies and provide guidance on how these should be addressed.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ High-Profile Role-Based Exceptions.

One of the most challenging role-based exceptions to address in a cybersecurity program is informally known as the “Executive Exception.”⁹⁸ It alludes to senior management not wanting to or not being able to follow the guidelines and controls of their organization. Such exceptions are a danger due to executives’ access to and authority over a great deal of sensitive information and essential systems. Sometimes the desired exceptions are a case of a simple misunderstanding or lack of understanding (such as an executive having access to systems that are unnecessary for their role). Alternatively, the organization’s leadership may want accommodations to facilitate their duties (*e.g.*, access to PII when working from home or traveling) beyond the measures put in place for a standard user.

The RCO should develop a shared understanding of the benefits of organizational policies and controls. Practical tactics to impart the importance of a policy or control include: to clearly describe why compliance is necessary; and to speak the language of the cost associated with organizational risks.⁹⁹ However, the RCO may still need to shape service solutions to meet the needs of senior leadership and other personnel desiring role-based exceptions.

⁹⁸ <https://www.helpnetsecurity.com/2020/05/28/make-security-exceptions/>.

⁹⁹ See the Roadmap section “Follow” in **Must 9 (Policy)**.



Must 9: Policy

Organizations must develop, adopt, explain, follow, enforce, and revise cybersecurity policies.

“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.¹⁰⁰ 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.

Organizations must have cybersecurity policies and procedures to facilitate effective, efficient cybersecurity programs. Whether or not research cyberinfrastructure operators (RCOs) have a great deal of policy on other topics, the complexity and threats to mission success posed by cybersecurity mean that some minimum set of cybersecurity policies and procedures is a necessity.

Policies have a lifecycle. Ensuring that each phase of the policy lifecycle is implemented is key to an organization’s ability to make effective use of policy:

1. **Development** is the process of going from a “blank page” to documented statements suitable for adoption and use. These can be based heavily on templates and examples, or born from the experience and mission needs of the organization.
2. **Adoption** is the formal step where the organization recognizes the policy as one it will use and enforce.
3. **Explanation** is the process of giving the relevant people notice of and access to policy, and engaging in activities designed to help those people understand it.
4. **Following** is the process where the relevant actors carry out the policy’s requirements. The need to follow policy is as important as it is obvious.
5. **Enforcement** is the process of responding to behaviors that stray from the directives laid out in policy. The term has very serious connotations, but cybersecurity policy enforcement can include activities as gentle as friendly reminders.
6. **Revision** is the process of changing (and sometimes replacing or retiring) policies so that they remain in line with mission needs.

Contemporary organizations typically formalize cybersecurity policy as a collection of separately developed and adopted artifacts, rather than creating a single, monolithic “Cybersecurity Policy.” Trusted CI recommends this approach because it aids the drafting and revision of singular, topical policies over time.

There are many types of cybersecurity and cyber-relevant policies that RCOs may need, including

¹⁰⁰ Some organizations draw explicit distinctions between policy, procedure, requirements, specifications, and guidance. Hard distinctions may be operationally useful or necessary. The Trusted CI Framework does not draw hard distinctions because this Must and related guidance are relevant to all authoritative documented statements meant to guide action.

but not limited to: acceptable use, access control, asset management, disaster recovery, incident response, training and awareness, password length and complexity, physical security, social media, and vulnerability management. Because there are so many potential policies to develop, it makes sense to ask “what should my RCO prioritize?”

Trusted CI recommends that RCOs at a minimum prioritize 1) a master policy that defines the major structures and governance functions of the cybersecurity program, 2) an incident response policy and procedural guidance, and 3) acceptable use or related policy/agreement that sets expectations for the researchers’ use of the RCO’s information assets.

This small collection of policies may well be insufficient, but they are absolutely essential. RCOs may find that additional policies warrant high prioritization.

Why is this a Must?

Why have policies? An organization cannot enforce a policy that does not exist. While all organizations will have norms established by other means, informal norms are difficult to monitor, explain, change, and (particularly) enforce effectively or consistently.

Without policy, norms can evolve and decisions can be made and communicated, but organizations of all kinds reach a point in their evolution and lifespan when formalized policy adds predictability, efficiency, and structure needed to achieve the organization’s goals.

Why have a policy lifecycle? Policies are only worth the effort if they are carried through the activities of development, adoption, explanation, following, enforcement, and revision: *i.e.*, the policy lifecycle. Without adopting a policy, the result of development is merely text on a page and is not enforceable. If a policy is developed and adopted, but never shared with the relevant actors and never explained (*e.g.*, as the subject of training), it cannot be effective as guidance to action and it cannot be enforced fairly. A policy cannot have its intended effect unless it is followed by those to whom it applies. Without enforcement, a policy holds no weight in an organization. Finally, policy revision reflects the fact that missions, obligations, technologies, and the environment all change, often without organizational intent driving those changes: organizational policies do not automatically change with them. Moreover, even without these changes, policy may need revision for completeness and clarity.

The Roadmap

This section describes the steps needed for an RCO to develop, adopt, explain, follow, enforce, and revise cybersecurity policies. There are few hard gates in the policy lifecycle: if you finish a phase of the cycle you are probably already doing things in the next phase.

→ Step 0. Take stock of the policy environment.

Before beginning to develop a new policy, consider the following:

- a) The RCO's existing policies (whether cybersecurity-related or not) and how they affect the cybersecurity program and the policies you plan to develop. For instance, are any of these existing policies so out of date that they will need to be retired or revised? Is there an established approach for developing and adopting policy?
- b) RCOs often exist under the umbrella of parent organization's policies and collaborate closely with peer organizations. These organizational relationships both may provide opportunities and pose challenges at different stages of the Policy Lifecycle. Parent or peer organizations can be a source for established norms and policies which can be inherited or shared. In some cases, this can largely negate the need for an RCO to develop and maintain its own versions of those specific policies. For example, if the RCO's parent organization requires two-factor authentication via policy and the policy applies directly to the RCO, then there is no need for the RCO to develop and adopt a policy that requires two-factor authentication. Instead, the RCO would be better served by focusing on its role in explaining, following, and enforcing the parent organization's policy.
- c) Expectations of stakeholders and obligations to third parties as outlined in **Must 2 (Stakeholders & Obligations)**.
- d) Baseline controls (*see* **Must 15 (Baseline Control Set)**) that call for specific policies and procedures.

→ Step 1. Develop.

Policies to prioritize

There are many types of cybersecurity and cyber-relevant policies that RCOs may need. We recommend an intentional approach to selecting what normative guidance warrants formalization and determining which ones to prioritize. We discourage adopting policies or provisions that present personnel and stakeholders with impractical requirements, or that the RCO has no intention or capability to enforce.

If in doubt, Trusted CI recommends that RCOs prioritize:

1. **A master information security policy.** This describes the fundamental operational details and commitments of the cybersecurity program. Without this, an RCO is missing the fundamental formalization necessary for a cybersecurity program.
2. **Incident response policy and procedural guidance.** Because of the high tempo, high stakes nature of some cybersecurity incidents — and the fact that incidents can and do happen — having and exercising these policies and procedures is critical before “things go boom.”
3. **Acceptable use or related policies and agreements.** These communicate expectations for stakeholders (*e.g.*, researchers) who use the RCO's information assets. If crafting an acceptable use policy for a multi-institution collaboration or federation, consider the

acceptable use policy developed by WISE.¹⁰¹

Once those policies are in place, determinations as to what other policies to develop should be made based on whether there is a need. Policy needs can be driven by stakeholder expectations and obligations (*see* **Must 2 (Stakeholders & Obligations)**) as well as cases where the RCO identifies types of action or behaviors for which less formal means of encouragement or enforcement prove insufficient.

🔗 Trusted CI maintains templates for **Master Information Security Policy & Procedures**, **Incident Response Policy**, and several others at <https://trustedci.org/framework>.

If an RCO identifies that multiple policies are needed, consider staggering the start of development. During the process of developing a policy there will be lessons learned which should be incorporated into subsequent efforts. These lessons learned can be codified into a formal protocol for developing policies and these can be the basis for future policy development efforts.

🔗 For RCOs without a formalized policy for developing policies, Trusted CI maintains a **Policy Development Protocol Template**.¹⁰²

Roles

In general, there are three key roles to consider when identifying who should be involved in policy development efforts and in what capacity:

1. **Authors:** Responsible for writing, getting appropriate input on, and seeking adoption of a policy. We encourage policy writing as a small team (*e.g.*, 2-4 people) if possible.
2. **Reviewers / Input Providers:** Provide the primary source of input and feedback for the authors of the policy. These can be experts in a particular topic addressed by the policy, personnel and stakeholders (*e.g.*, researchers) who will be affected by the policy, or other key stakeholders. We strongly encourage seeking input from at least a sampling of the classes of people who will be expected to follow or implement the policy. This provides an opportunity to “sanity check” the clarity and feasibility of the policy, and increases the likelihood that people will follow the policy when adopted.
3. **Approvers:** Have authority to adopt the policy. These can be individuals (*e.g.*, the RCO’s director or specific domain authorities such as risk officers, human resources, or legal counsel), or groups (*e.g.*, governing boards).

RCOs should anticipate the need to actively manage interactions with input providers in order to get sufficient feedback. Although in some situations the set of authors can deliberately include specific subject matter experts who have sound input to codify into policy, this should certainly not be the sole source of information when developing policy. High-quality input and feedback from those to

¹⁰¹ <https://wise-community.org/wise-baseline-aup/>

¹⁰² Check out Trusted CI’s **Policy Development Protocol Template**, available at <https://trustedci.org/framework>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

whom the policy applies is the most effective method authors have to acquire the information they will need in order to develop an effective policy.

High-quality feedback can often require active engagement over multiple revisions. To achieve this efficiently, authors should scope their feedback-seeking interactions appropriately for the audience. Different policies—and even different sections of a single policy—will be relevant to different input providers. Seek engagement from specific stakeholder groups on those items which are most applicable to them. Because of the scale of these interactions, it will be very important for authors to keep records of key decisions and conversations with input providers throughout the process so that important determinations are not lost.

Common Provisions

All policies should have an adoption date/version and signature or other marking to show official adoptions, as well as provisions for enforcement and attaining exceptions.

When developing enforcement provisions, RCOs should be sure to address the following:

1. What is the process and who is responsible for explaining, monitoring/verifying adherence, and enforcement?
2. Is there an escalation process? To whom? Legal? Human Resources? Law Enforcement? Who would make that determination?
3. Describe any conditions which would lead to an automatic action (*e.g.*, disabling an account).

Likewise, when developing exceptions provisions, RCOs should address:

1. What is the process and who is responsible for determining whether exceptions are granted?
2. Do they actually have the authority to accept any risk introduced by the exception?
3. How are approved exceptions documented and tracked (particularly if temporary)?

Templates and Consistency

Rather than “reinventing the wheel,” use templates and examples from peer organizations to get started on new policies. Organizations such as SANS,¹⁰³ AARC,¹⁰⁴ and Trusted CI¹⁰⁵ all have policy templates publicly available.

In order to be able to maintain a consistent policy base, RCOs should consider incorporating common themes into RCO policy templates and using these templates as a starting point for all policies. Starting with a sound template as the base provides several benefits when developing policy.

1. It ensures that crucial components such as exception and enforcement provisions are included by default.
2. It ensures that definitions and terminology used throughout an organization’s policy base are

¹⁰³ <https://www.sans.org/information-security-policy/>.

¹⁰⁴ <https://aarc-project.eu/policies/policy-development-kit/>.

¹⁰⁵ <https://www.trustedci.org/framework>.

consistent.

3. It sets formatting and style standards in a manner consistent with other policies.

→ Step 2. Adopt.

By adopting a policy, an organization begins the commitment to explain, follow, and enforce the policy. Once a policy is formally adopted (*e.g.*, by senior RCO management), it is someone's responsibility to ensure that the policy is shared, explained, followed, and enforced. RCOs should be wary of adopting policies when these responsibilities are not spelled out explicitly. RCOs should have a formal repeatable procedure for policy adoption. Again, policies need an adoption date/version and signature or other marking to signify the adoption.

→ Step 3. Explain.

Explanation encompasses the processes of giving the relevant people notice of and access to policy, as well as engaging in activities designed to help those people understand it.

Notice

It is unreasonable to expect people to follow a policy if they don't know of its existence and contents. Policies should be made readily available and stakeholders should be reminded where they can find policies. Certain compliance regimes such as HIPAA require this explicitly. For some types of policy it may be appropriate to require acknowledgement of receipt or reading of the policy. Acceptable Use Policies for systems or resources made available to third parties, for example, may be more readily enforced when those third parties have to acknowledge their receipt of the policies to which they must adhere before they can use these resources.

Early Engagement

In some cases, policy development is where this process begins. Seeking input from the right people introduces the policy's key ideas early in the Policy Lifecycle. In ideal situations, individuals who are engaged in providing input and feedback become advocates for the policy within the organization.

Training

Formal training is another common means by which policy is explained. Training can be simply made available or can be a requirement of the policy itself. If training is the primary means of explaining a policy consider the following:

1. How will this training be carried out? In person? Is it pre-recorded?
2. When and how often should the training be completed by individuals?
3. How will the organization track completion and effectiveness of the training? Do trainees acknowledge completion? Is there any testing? Are the results recorded?
4. What resources will be needed to provide the training? Who will provide them?
5. When and how will the training be revised?

Other Means of Explanation

Formal training is not the only approach an organization can adopt by which to explain policy. For example, tabletop exercises can provide highly interactive simulations in which to explore the implications of policies. These interactive exercises can be significantly more engaging than generalized training, but can also require significantly more resources and expertise to carry out effectively. Moreover, instances of policy enforcement can lead to opportunities to explain and clarify policies.

→ Step 4. Follow.

The need to follow policy is as important as it is obvious. Once a policy is adopted and explained, the relevant actors need to carry out the policy's requirements.

→ Step 5. Enforce.

Enforcement comes into play when an RCO identifies that a person or group is not following the policy. While enforcement provisions need to be clear about the potential consequences of violation (*e.g.*, loss of access; loss of employment; reporting to law enforcement), enforcement is about more than people getting in trouble: it includes simple reminders, putting controls in place to prevent violations, and correcting honest mistakes. The strategies under Step 3: Explain can be the most efficient, effective means of policy enforcement and reinforcement.

A policy that is never enforced is at best a suggestion, and could be a source of cybersecurity vulnerability or liability.¹⁰⁶ Moreover, policies should be enforced consistently according to the enforcement provisions of the policy itself. If the policy cannot be enforced consistently or effectively, RCOs should consider retiring or revising the policy.

RCOs will not achieve perfect adherence to policies at all times. In order to both enforce the policy and discover cases where the policy is ineffective, adherence to the policy should be monitored. RCOs may discover not only noncompliance that needs to be corrected, but also edge cases where exceptions could have been granted or places where policy is keeping people from being effective at their jobs. RCOs should make note of these situations when they are discovered as they may indicate that a revision to the policy or some other change is appropriate. Consider the following:

1. Who is not following? Does it actually matter whether they follow the policy? If not, this may indicate that a revision is appropriate in order to more carefully specify whom the policy applies to.
2. Do they actually have good reason not to follow? If so, a revision to the exception procedures in the policy may be appropriate.
3. Do they know about the policy? If not, it may be appropriate to reconsider the strategies for explaining the policy.

¹⁰⁶ *E.g.*, in the context of cybersecurity and privacy compliance regimes (*e.g.*, HIPAA), it is commonly understood that regulators view failure to enforce policies even more negatively than a lack of policy.

→ Step 6. Revise.

As with the cybersecurity program generally, policies need to be refined, and sometimes retired or replaced altogether. The reasons for revising policies will be situationally dependent. Some policies may be stable for years; others may need frequent updates driven by experience or environmental changes. Incident response policies for example can benefit greatly from a revision cycle when taking advantage of the lessons learned throughout the response process. Consider the following:

1. Periodically review all cybersecurity policies. This could be done as part of the comprehensive programmatic review contemplated in **Must 10 (Evaluation & Refinement)**.
2. Mark superseded policies and archive them with a reference to the current version.
3. Track policy revisions. For example, implement a timestamped change log within the policy itself so people can easily find what changed and when.

Common Challenges & Recommendations

This section describes a common **Must 9** challenge and offers recommendations on how to overcome it.

→ When “policy” is a trigger word.

For some organizations, “policy” denotes a highly formal statement of the organization and implicates heavyweight processes for the Policy Lifecycle. In some cases, an RCO may need and want to take cybersecurity policy through these rigors. For instance, a master information security policy—with its high-level statements of how cybersecurity will be governed—may warrant the full attention and ultimate force of heavier adoption processes. However, RCOs should keep in mind that—for the purposes of **Must 9**—policy refers to **any documented normative statements adopted by an organization to govern human behavior**.

For the purposes of the Framework, these include authoritative documented statements of “policy,” but can also include “procedures,” and other normative guidance. For RCOs (or RCO parent organizations) with distinct tracks for “policy” and “procedure,” “guidance,” or “standards,” consider the implications of the labels and whether the desired outcome can be achieved without engaging the most time-consuming or labor-intensive path. Unlike a master policy, vulnerability management or password complexity decisions may not require the same attention and status to achieve their ends effectively, efficiently, and fairly.



Must 10: Evaluation & Refinement

Organizations must evaluate and refine their cybersecurity programs.

Programmatic evaluations are how the organization determines whether the cybersecurity program is achieving its purpose. Refinements are any changes designed to improve the program's efficiency or effectiveness. Evaluation and refinement of a cybersecurity program can take many forms depending on the formality and scope of the assessment and the type of evaluation (*e.g.*, planned, comprehensive program evaluations; internal self-evaluations following an incident).

Evaluations happen at multiple levels. Organizations should undertake preplanned, periodic “comprehensive” evaluations, where the efficacy of the entire cybersecurity program is considered and where the strategic plan (created in **Must 1 (Mission Focus)**) is reaffirmed or modified. However, other evaluations conducted in response to changes or events can impact the program. These “trigger events”¹⁰⁷ are typically limited to a portion of the program on a particular issue and offer an opportunity to assess and refine tasks or activities found in individual Musts' roadmaps.¹⁰⁸ Additionally, these smaller, event-driven evaluations form an information base that allows the program's comprehensive evaluation to be more successful.

Evaluations can take many forms, with different types of evaluations offering different strengths and weaknesses.¹⁰⁹ The most important distinction is between evaluations conducted internally and evaluations carried out by an external organization. External evaluations are more costly and time-consuming, but offer a more objective perspective and allow external experts to provide valuable input into the cybersecurity program's workings.

Trusted CI recommends a comprehensive external assessment of the cybersecurity program¹¹⁰ every three to five years.¹¹¹

The comprehensive evaluation should be informed by this Framework,¹¹² the mission, and the

¹⁰⁷ A trigger event can be technical (*e.g.*, firewall modification or network configuration changes, IdM access-group definition changes, local administrator installing new remote access software) or organizational (*e.g.*, shifts in unit service responsibilities, new or modified service agreements with a third party, new guidelines or requirements implemented by sponsoring entities). The technical changes might be ordinary change management if the expected impact is minor.

¹⁰⁸ An RCO should develop its own list of these events that “trigger” an evaluation. Examples of events that would trigger an evaluation for Musts are included in the Appendix B.

¹⁰⁹ The Cybersecurity Assessment Parameter Profile (CAPP) was developed as a tool for decision makers to: “(1) identify the salient differences between existing cybersecurity assessments; (2) select the most appropriate cybersecurity assessments for their missions, resources, and constraints; and (3) find and fill gaps in the cybersecurity assessment ecosystem.” See <https://apps.dtic.mil/sti/citations/AD1049123>.

¹¹⁰ While an internal evaluation is possible, see “Selecting the External Review Team” in the Common Challenges and Recommendations section concerning the types of internal and external reviews and considerations thereof.

¹¹¹ RCOs dealing with certain classifications of regulated data may have a requirement for a periodic assessment that is more frequent.

¹¹² If the RCO is unable to find a commercial assessor to conduct this evaluation, peer organizations may be more acquainted with the Framework and able to serve as external assessors.

cybersecurity strategic plan developed in **Must 1 (Mission Focus)**. The evaluation scope includes the past, present, and future detailed operational and tactical plans for the cybersecurity program.

Why is this a Must?

Both the organization and the organization's environment change, and the cybersecurity program needs to change with it. The organization must take a step back and consider whether these changes impact the program and look for ways to improve the program's alignment. Comprehensive evaluations ensure the program continues to track with the Framework, the organization's missions, and its cybersecurity strategic plan.

Organizations also experience incidents and other unexpected cybersecurity-relevant events that may require more immediate action. In the aftermath of these events, organizations need to evaluate and refine the cybersecurity program based on the lessons learned.

Finally, cybersecurity programs are complex and challenging to get right, and organizations will need to iterate and make refinements based on lessons learned. Evaluations provide an opportunity to realign resources, re-prioritize activities, and inform decision-making about the program by research cyberinfrastructure operator (RCO) leadership.

The Roadmap

This section describes the steps needed for an RCO to evaluate and refine their cybersecurity program. Implementing **Must 10** involves two related activities: a periodic, comprehensive program evaluation; and evaluations in response to "trigger events". This roadmap lays out a process for formalizing both the comprehensive evaluation and the ongoing trigger event evaluations. Step 1 is to plan for a comprehensive evaluation. Step 2 is to define a standardized process to document ongoing changes in between comprehensive evaluations. Step 3 is to execute the plan for the evaluation. Step 4 is to refine the program and adjust the plan for future evaluations.

→ Step 1. Plan for a comprehensive evaluation.

The goal of this step is to formalize an organizational commitment to engage in program evaluation and revision. This formalization includes documenting: who will conduct the evaluation (*i.e.*, an internal or external¹¹³ evaluator);¹¹⁴ the relevant roles and responsibilities;¹¹⁵ any delegation of evaluation activities; and the types of expected outcomes.

Ideally, the requirement to conduct a comprehensive evaluation is documented in the organization's

¹¹³ For a discussion of internal vs. external evaluations, see Russell & Jackson, "Cybersecurity Assessment Parameter Profile (CAPP): A Tool for Making Sense of Cybersecurity Assessments," March 2018, <https://apps.dtic.mil/sti/pdfs/AD1049123.pdf>.

¹¹⁴ See "Selecting the External Review Team" in the Common Challenges and Recommendations section.

¹¹⁵ The roles and responsibilities are particularly essential to keep updated since personnel changes are likely over time, and some activities (*e.g.*, data collection) are best performed on an ongoing basis.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

Master Information Security Policy & Procedures (MISPP)¹¹⁶ or the organization’s equivalent of a MISPP (*see* **Must 9 (Policy)**). Additionally, this step may include formally documenting the requirement to monitor and catalog “trigger events” that impact the organization’s cybersecurity program.

→ Step 2. Define a Standardized Process for Documenting Changes.

Since the comprehensive program evaluation occurs infrequently, there will be significant periods between evaluations. However, this does not preclude evaluation and refinements during this time. The RCO will continue to refine individual Musts and respond to emergent “trigger events” that impact the program. This step aims to ensure that the RCO documents any changes made to the cybersecurity program in response to these smaller evaluations and refinements. The documentation helps ensure the RCO has relevant information needed when conducting its comprehensive evaluation.

Documentation should include, at a minimum, a preliminary discussion of what types of events will qualify as trigger events and how the RCO should catalog the trigger event and the actions taken in response. There is a benefit to standardizing how the RCO documents trigger events. Having a standardized process for documenting changes streamlines the process and helps ensure that changes are documented consistently and promptly.

The RCO should establish clear communication lines between organizational units with respect to the cybersecurity program. Robust communication within the RCO helps avoid undetected changes that could impact the program. Fostering a sense of cooperation and shared expectations can facilitate awareness of governance changes, shifts in stakeholder expectations, and modifications to information systems or services.

→ Step 3. Execute the comprehensive evaluation.

The comprehensive evaluation can be conducted by either an internal or an external team, but Trusted CI recommends that RCOs use an external team whenever possible. An external assessment’s objective analysis can uncover systemic problems, gaps, and inconsistencies in implementation that RCOs may not have identified in a self-assessment. Additionally, an external assessment can provide information comparing the current cybersecurity program with that of comparable organizations.

With any review, it is important to have an up-front agreement on the engagement plan. All parties need to understand the expected reviewer roles and responsibilities, what is being evaluated, the standard to be evaluated against, the timeline, the expected resources required, and the development, review, and approval process for the final report.

For more on selecting external assessors, *see* Common Challenges and Recommendation “Selecting the External Review Team.”

¹¹⁶ Check out Trusted CI’s **Master Information Security Policy & Procedures Template**, available at <https://trustedci.org/framework>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ Step 4. Refine the program and adjust the plan.

The recommendations and lessons learned from the comprehensive evaluation are now incorporated into the cybersecurity program. These refinements may include changes to the cybersecurity strategic plan, changes to individual **Must** implementations (*e.g.*, how the RCO allocates resources or how it documents risk acceptance), and changes to other parts of the organization. Although all recommendations do not need to be followed blindly, the RCO (and its parent organization, where applicable) should thoughtfully consider them. The RCO should document the rationale when not taking action on significant recommendations.

Additionally, reviewing the process of undergoing a comprehensive evaluation will help with future evaluations. The RCO should document potential areas for improvement or lessons learned for future evaluations. For instance, if there was information requested by the evaluators that the RCO was unable to provide or was challenging to obtain, the RCO should work to make that information more easily accessible in the future. Moreover, the RCO should look for areas that would benefit from more rigorous monitoring and review between comprehensive evaluations. Finally, take time to celebrate what went well, especially if the evaluation uncovered significant issues.

Common Challenges & Recommendations

This section describes a common **Must 10** challenge and offers a recommendation on how to overcome it.

→ Selecting the External Review Team.

An RCO that decides to utilize an external team for its comprehensive evaluation has many options to consider when making its decision. An *ad hoc* group can be assembled from a diverse collection of internal stakeholders external to the cybersecurity team. In some cases, the RCO or parent organization may wish to select stakeholders external to the RCO to participate in the assessment. Note that RCOs may have stakeholder constraints or external obligations that necessitate a specific type of evaluation or particular evaluator. For instance, for larger RCOs, there may be explicit requirements as to which party conducts a cybersecurity assessment. Similarly, RCOs dealing with categories of protected information (*e.g.*, CUI, ePHI) may be required to have a periodic review by a commercial auditing firm.¹¹⁷

Peer organizations can provide added objectivity with a fresh perspective on risks and recommendations. Reviews by peer organizations can provide invaluable insight into a program's maturity. There is an added benefit for the participating reviewers who may see solutions the RCO has implemented to apply to their peer organization's problems.

Similarly, industry-specific organizations such as Trusted CI or REN-ISAC perform external, third-party reviews. These organizations may have a broader perspective than peer organizations and have a great deal of experience performing comprehensive reviews of cybersecurity programs.

¹¹⁷ Consulting peers with the same compliance challenges may provide a list of firms with experience assessing research and academia.

Finally, RCOs can always select from a wide swath of commercial assessment companies. When selecting a commercial offering, it is critical to communicate the assessment's scope and purpose. This communication avoids receiving a more conventional penetration test, risk assessment, or other evaluation types not scoped to the cybersecurity program.



Resources



Must 11: Adequate Resources

Organizations must devote adequate resources to address unacceptable cybersecurity risk.

The organization's cybersecurity program requires resources to protect the organization's mission. These include budgeted funds and personnel, as well as external resources (*e.g.*, cybersecurity tools and services). An adequate level of each type of resource must be dedicated to cybersecurity. If an organization determines that the magnitude of cybersecurity risk is unacceptable, then resources must be brought to bear to address that risk.

Cybersecurity programs require money and personnel to be effective, and require additional resources like tools and services to efficiently carry out their cybersecurity activities. The allocation of each type of resource must be adequate for the program's needs. **Must 11** sets the standard for the remaining Resources Musts,¹¹⁸ answering the question "how much?"

"Adequacy" is ultimately a determination made by organizational leadership to establish that the current investment in cybersecurity is sufficient and that outstanding cybersecurity risks facing the organization's mission are acceptable.¹¹⁹ What is adequate for one organization may be inadequate for another. These are decisions for organizational leadership to make, taking into account the organization's mission, risk, and willingness to accept risk. These decisions should be strategic, pragmatic, and realistic.

Adequacy is directly tied to **Must 6 (Risk Acceptance)**. This is the "put your money where your mouth is" Must. If a risk is determined to be unacceptable, more resources are needed to address that risk. Conversely, if the organization decides that the current investments in cybersecurity are adequate, any outstanding risks can be considered *de facto* accepted. Each time resource allocation is reassessed (*e.g.*, during annual staffing and budgetary planning and approval), leadership will need to consider the organization's current cybersecurity risk,¹²⁰ the organization's planned resource allocation, and make a determination that those allocations are adequate.

Moreover, adequacy is not a one-time determination. Risks that are unacceptable over the long term may be acceptable in the short term, and organizations may decide that their current resource allocation is adequate despite outstanding long term risks. Indeed, frequently organizations will defer action on particular controls due to budgetary constraints, with the full intent to address those controls in future budgetary cycles (when more resources are available or current priorities are completed).

Finally, it is important to note that resources in one category may necessitate resources in other

¹¹⁸ The other Resources Musts include **Must 12 (Budget)**, **Must 13 (Personnel)** and **Must 14 (External Resources)**.

¹¹⁹ Note, "adequacy" as used here is not a legal conclusion or related to a third party standard: it is an internal determination made by the organization.

¹²⁰ An organization's overall cybersecurity risk is difficult to determine, and a full discussion is outside the scope of this Must. For a discussion of inputs to risk acceptance decision making, *see* **Must 6 (Risk Acceptance)**, Common Challenge and Recommendation "Inputs to risk acceptance decision making".

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

categories. For instance, many tools and services require personnel time to manage, and so the investment in a new tool may also require a commensurate investment in personnel hours to install, configure, operate, and maintain the tool. Absent both investments, the new tool will simply become “shelfware,” and provide no tangible benefit to the organization. For more discussion of external resources, *see* **Must 14 (External Resources)**.

Why is this a Must?

Organizations need to devote resources to cybersecurity because, fundamentally, cybersecurity isn’t free. It requires money, it requires people, and it requires support from the organization. It is not enough to allocate token or minimal resources to cybersecurity: the resources must be adequate to address the risks the organization faces.

Many cybersecurity frameworks forego any discussion of resourcing cybersecurity. Instead, they focus purely on lists of security controls, without discussion of the organizational realities required to actually implement those controls. It is important that the organization recognizes that when it commits to implement a cybersecurity program, it is also committing to adequately resource that program, just as it would with any other business function.

Finally, adequacy is not designed to be a high bar. An organization may allocate very few resources to cybersecurity, but still be “adequate” so long as their leadership consciously accepts the risk from such low resourcing. The goal of adequacy is to make sure that leadership recognizes that if they have “unacceptable risks” that they do not provide resources to address, then they do not have adequate resources allocated to cybersecurity.

The Roadmap

This section describes the steps for a research cyberinfrastructure operator (RCO) to determine if it has allocated adequate resources to its cybersecurity program. The basic process is: Step 1, understand your current state; Step 2, have leadership determine adequacy; Step 3, develop a plan of action; and Step 4, implement and evaluate the plan.

The following steps should not be understood as requiring a monolithic, standalone process. Rather, these are the steps the RCO can undertake whenever making a decision regarding the adequacy of their cybersecurity resources, such as during annual budgeting, personnel turnover, or in response to an incident or evaluation/assessment.

→ Step 1. Understand Your Current State.

The first step is to understand the current state of your organization, what resources you currently allocate to cybersecurity, and what “adequate” means for your organization.

Your current state is a complex topic, but will definitely include understanding your implementations of: **Must 1 (Mission Focus)**; **Must 2 (Stakeholders & Obligations)**; **Must 4 (Asset Classification)**; **Must 11 (Adequate Resources)**; **Must 15 (Baseline Control Set)**, and **Must 16**

(Additional & Alternate Controls). Additionally, your current state must consider the risks your organization faces, and how well the organization has mitigated those risks.

Moreover, understanding the organization's current state requires understanding what "adequate" means for that organization. Organizations that are more comfortable accepting greater risks will define "adequate" differently than those that are less comfortable accepting risks. Because cyber risk is almost impossible to quantify, this "risk tolerance" should be understood as a reflection of how willing or unwilling organizational decision makers are to accept risk.

→ Step 2. Have Leadership Determine "Adequacy".¹²¹

The second step is to bring together organizational leadership to determine if the allocation of resources for the cybersecurity program are adequate.¹²² The goal with this step is to ensure that the organization's leadership is apprised of the current state of the organization's cybersecurity and is able to make an informed decision¹²³ about whether the current resources are adequate or if the organization needs to: 1) allocate greater funds; 2) hire additional personnel; and/or 3) purchase additional tools/services.¹²⁴

Additionally, this is where the organization's leadership should identify any problems that indicate the current resource allocation is inadequate. The most direct indicator is that the organization has risks it has deemed unacceptable, but has no plan to allocate additional resources to mitigate those risks (at least in the foreseeable future). However, other common signs of problems include: 1) Cybersecurity tools that were purchased but are not being utilized; 2) individuals performing cybersecurity duties despite not having security in their job description;¹²⁵ 3) leadership being unaware of or willfully blind of cybersecurity risks; and 4) cybersecurity incidents.¹²⁶

→ Step 3. Develop a Plan of Action.

The third step is to develop a plan to allocate the resources approved by the organization's leadership and lay out a longer term strategy that maps future resources to planned future mitigations. Since organizations will rarely implement all of the desired mitigations in a single budgetary cycle, leadership will need to accept the residual risk from not implementing particular controls until future budgetary cycles. Note, this plan of action may be included within the Cybersecurity Program Strategic Plan discussed in **Must 1 (Mission Focus)**.

For more details on how to prioritize limited resources, *see* the Common Challenge &

¹²¹ Note, Steps 2 and 3 may be reversed in some organizations. In such a case, Step 2 would be to Develop a Potential Plan of Action, and Step 3 would be for Leadership to approve, reject, or modify the plan.

¹²² Adequacy determinations are complex decisions, where RCO leadership will consider a number of factors, including but not limited to: regulations, contractual obligations, security incidents, and adherence to common best practices.

¹²³ Major decisions regarding organizational resourcing should be documented so they are reviewable after the fact.

¹²⁴ Note, this step is also where leadership would determine if current resources are above and beyond adequate, and if the organization should commit additional resources above what it has deemed "adequate."

¹²⁵ Both items 1 and 2 may indicate a lack of personnel resources.

¹²⁶ Note, however, that suffering a cybersecurity incident does not necessarily mean that resources were inadequate prior to the incident. However, it is fairly common for organizations to reassess what is "adequate" after experiencing an incident, and prioritize cybersecurity higher than before the incident.

Recommendation “Do I spend money I don’t have, or accept “unacceptable” risks?”

→ Step 4. Implement and Evaluate the Plan.

The final step is to implement the plan generated in Step 3 and evaluate its effectiveness. As resourcing is an ongoing challenge, and plans often don’t survive first contact with reality, organizations will want to periodically evaluate their resourcing to determine if what was initially deemed “adequate” is adequate in practice. Additionally, exactly what constitutes “adequate” for your organization may change based on changes to your mission, threats, risks, incidents, stakeholder environment, current leadership, and a host of other factors.

Common Challenges & Recommendations

This section describes some common **Must 11** challenges and offers recommendations on how to overcome them.

→ But how much should I spend?

Although adequacy necessarily is tied to risk acceptance, in practice, risk acceptance is a messy and complicated topic. Organizations may not feel comfortable relying solely on their internal risk acceptance decisions to determine if it is spending enough money on cybersecurity, and may prefer a more concrete standard. In these cases, organizational leadership can choose to focus on setting an adequate total cybersecurity budget, and working backwards from there.

Some sources of external guidance related to cybersecurity budget are: benchmarking data, expert recommendations, as well as comparisons with one’s peers. (Note, benchmarking data typically will show lower budget numbers than expert recommendations, as benchmarking data simply reflects what other organizations are spending, and does not mean that those organizations’ spending is itself adequate.) For instance, a common expert recommendation for cybersecurity budgets is between 8% and 10% of IT budget.^{127,128} Alternatively, benchmarking research typically finds that organizations spend between 3% and 12% of IT budget on cybersecurity.¹²⁹ Note, however, that both benchmarking research and expert opinions typically emphasize that factors like size and business sector play an important role in setting budgets, with, for instance, smaller organizations requiring proportionately larger budgets than larger organizations.¹³⁰

Much external guidance will be general (*i.e.*, not tailored to the specific needs of the organization) and will require human judgment. Organizations can increase the relevance of this information by looking at materials specifically targeting their business sector, organizations of similar size, and organizations with similar threat profiles. Organizations can also enlist the help of third party assessments to look closely at the organization and their missions.

¹²⁷ See, e.g., cybersecurity budget recommendations of Richard A. Clarke in “The Fifth Domain,” <https://www.npr.org/2019/07/16/742386872/a-look-at-the-vulnerabilities-and-capabilities-of-american-cybersecurity>.

¹²⁸ See also related discussion in **Must 12 (Budget)**, “Benchmark Your Budget.”

¹²⁹ See <https://scholarworks.iu.edu/dspace/handle/2022/22289>.

¹³⁰ *Ibid.*

Based on this external guidance, the organization can set its budget. Although the external guidance does not set hard boundaries on what the budget can be, decisions to go outside of the established range should be done with appropriate consideration. If an incident occurs and the organization's budget is below the established norm, stakeholders will likely ask difficult questions about why the budget was so low. For a more complete discussion of cybersecurity budgets, *see* **Must 12 (Budget)**.

→ Where do I find the resources?

A natural concern for organizations is where they will find more internal resources to allocate to cybersecurity. Organizations rarely have unallocated resources ready and waiting, and getting more resources for cybersecurity (commonly viewed as a cost-center rather than a value generator) can be particularly challenging. Moreover, RCOs often lack options typical in other communities like taking on debt, lowering profits, and raising prices.

In general, there are three areas where resources can be pulled from when spending cannot be unilaterally increased across the organization: 1) reallocation of internal resources; 2) external resources; and 3) free resources.

Internal reallocation is a natural place to start, but may also be the most difficult politically and practically: a cybersecurity budget increase means a budget decrease in another part of the organization. However, it is important to note that in organizations that under-allocate to cybersecurity, it is often the case that cybersecurity activities are still being performed, they simply are not reflected in the budget.¹⁵¹ This does not reflect “free work”; rather, it means that organizations are unaware of how much money they are currently spending on cybersecurity, and how much productivity is lost by personnel performing duties they were not hired (and likely not trained) for.

External resources¹⁵² is the largest category. These resources can range from private sector cybersecurity firms that can perform work more efficiently, to public sector agencies that offer grants or other aid for improving cybersecurity,¹⁵³ to the RCO's funding agencies. The role of funding agencies in resourcing cybersecurity is not settled, and so it is valuable for RCOs to communicate their needs to their funding agency to seek additional resources specifically for cybersecurity. A particularly important “external” resource is the organization's parent organization, which may offer a number of valuable resources at minimal cost.

Finally, free resources include options like open source libraries, public interest groups, and free online training. These “free” resources may offer less benefit than paid services, but come without a fee, and as such can be useful supplements, particularly when faced with common challenges for which others have developed adequate solutions. However, free resources should be considered carefully, as many “free” resources come with hidden costs (*e.g.*, required personnel to operate or maintain), or may have embedded vulnerabilities or inherent limitations.

¹⁵¹ *See* **Must 13 (Personnel)**.

¹⁵² *See* **Must 14 (External Resources)**.

¹⁵³ *See, e.g.*, <https://www.fema.gov/grants/guidance-tools/non-disaster-grants-management-system>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ Do I spend money I don't have, or accept “unacceptable” risks?

A potential challenge organizations may face is the prospect of a number of “unacceptable risks,” but a budgetary reality that precludes allocation of more resources to cybersecurity. This creates a Catch-22¹³⁴ where the lack of resources forces the organization to accept risks that it considers “unacceptable.” The best solution to this challenge is to ensure that current resources are allocated as efficiently as possible by effectively prioritizing what investments to make and in what order.

Prioritizing cybersecurity investments is a challenging topic. Outside of the cases (where there is a clear best answer, there is no universal model that RCOs can adopt to determine which investments to make. Instead RCOs will have to collect the best information they can, think critically about what their biggest security priorities are, and make strategic decisions.

When starting from scratch, prioritization for most organizations is fairly easy: prioritize the highest-value controls.¹³⁵ Controls such as application “allow lists,” two-factor authentication, offline backups, and patching are all known to be effective against a wide swath of attacks, and will provide the most “bang-for-your-buck.”¹³⁶ Similarly, a good baseline control set will be prioritized, and most organizations can follow this prioritization and trust that it will yield better results than attempting to reinvent the wheel or do an in-depth cost-benefit analysis for control selection.

Outside of these easy cases, however, prioritization becomes much more challenging. In general, organizations will want to consider four factors: 1) which cybersecurity outcomes the organization cares most about (*e.g.*, preventing system downtime; maintaining the confidentiality of privileged information);¹³⁷ 2) which threats are most prominent to the organization (*e.g.*, known ongoing phishing campaigns) and the organization’s business sector (*e.g.*, ransomware attacks in the organization’s business sector); 3) the RCO’s current controls; and 4) what activities the RCO’s leadership feels must be addressed promptly, and what activities leadership is willing to accept the risk from not implementing for the near term.

The first consideration is the RCO’s mission. When faced with limited resources, the RCO will need to make sure that those limited resources are targeted to maximize the value they provide to the RCO’s mission. *See also* **Must 1 (Mission Focus)**. It is valuable in this step to focus on the impact of risks, and determine what cyber outcomes would have the most negative impact on the organization, to help determine what the RCO cares most about protecting. The perspectives of the organization’s stakeholders (*e.g.*, funding agencies) will be valuable in making this determination.

The second factor is the threats the RCO is facing, both to itself and to similarly situated organizations. In general, RCOs should prioritize investments targeted at actual or imminent threats and avoid investments targeted at hypothetical or theoretical threats. RCOs can make use of freely available threat intelligence to inform their understanding of what threats they are facing.¹³⁸

¹³⁴ *See* <https://www.theidioms.com/catch-22/>.

¹³⁵ *See also* **Must 15 (Baseline Control Set)**.

¹³⁶ *See* <https://www.cyber.gov.au/acsc/view-all-content/essential-eight/essential-eight-explained>.

¹³⁷ Note, funding agencies can also play a significant role in determining what the organization “cares about.” When applicable, funding agency preferences should be included in this calculus.

¹³⁸ *See, e.g.*, <https://us-cert.cisa.gov/ncas>.

The third factor is the RCO's current control implementation. RCOs will want to consider areas where their cyber defenses are comparatively weak or assets that are particularly vulnerable. These areas may warrant increased resource allocation to shore up the current deficiencies.

The final factor is time. RCOs will often find that they have insufficient resources to address all of their actual and imminent threats in one budgetary cycle. In these scenarios, the RCO will need to pick and choose which investments to make this year and which to defer until a future budgetary cycle. RCOs therefore can be strategic with their resources, spacing their mitigation efforts over longer timescales, and establishing multi-year plans to address their unacceptable risks.

With these four factors in mind, RCOs should be empowered to make strategic decisions about how to prioritize their cybersecurity investments.¹³⁹

¹³⁹ Note: some frameworks attempt to exactly quantify the likelihood and impact of discrete risks and use these numbers to prioritize control selection. Although the motivation here is sound, the implementation is riddled with problems, and attempting to quantify risk in this way is unlikely to provide useful results. Having some understanding of the relative likelihood and impact of particular risks is valuable, but the numbers generated will never be accurate enough to establish a "correct" answer.



Must 12: Budget

Organizations must establish and maintain a cybersecurity budget.

Cybersecurity budgets are financial plans that commit specific resources for the organization's cybersecurity efforts over a designated period of time. Cybersecurity budgets serve an important documentation and planning function, clearly stating in advance which resources the organization is committing specifically to cybersecurity. Cybersecurity budgets allow for cybersecurity decision makers to plan and execute cybersecurity functions in a more deliberative manner, and demonstrate the organization's commitment to investing in cybersecurity.

Cybersecurity budgets may be standalone documents or may be part of a larger organizational budgeting process. The exact level of detail and granularity in a cybersecurity budget will vary based on the needs of the research cyberinfrastructure operator (RCO). Indeed, some RCOs may simply have a line item with the total budget for cybersecurity, while others may go into great detail breaking down how that budget should be distributed among specific cybersecurity functions.

This Must requires only that RCOs develop a financial plan that commits the organization to allocate funds to cybersecurity, along with the appropriate supporting documentation. It does not address the question of "how much" should be committed. Organizations with a cybersecurity budget can still be woefully under-allocated for cybersecurity or may spend those resources in ill-advised ways. The goal of the budget is to require leadership to think critically about what funds they are willing to commit to cybersecurity and to document their decision making on this issue in a manner that is transparent and reviewable after the fact.¹⁴⁰

Why is this a Must?

Must 12 is a technically simple requirement to satisfy, but it is a uniquely impactful way to improve organizational cybersecurity decision making. RCOs cannot make informed decisions about whether to spend more or less on cybersecurity if they do not know what they are already spending. The budgeting process provides cybersecurity professionals an opportunity to communicate with leadership about their needs and encourages those cybersecurity professionals to thoroughly justify why those "needs" are actually needed.¹⁴¹ This provides a recurring, structured process for leadership to seriously consider their cybersecurity posture, and whether they have devoted adequate resources to cybersecurity in light of the cyber risks and threats faced and their cyber risk tolerance. For a more complete discussion of this process, *see* **Must 11 (Adequate Resources)**. Absent an explicit budget, leadership can plausibly claim ignorance of its own cybersecurity posture, assuming

¹⁴⁰ Note: a significant factor impacting the content of an RCO's cybersecurity budget will likely be the role of any parent organization. The level of involvement and overlap between the RCO's cybersecurity program and its parent organization's cybersecurity program will impact how many resources are allocated in the budget. Essentially, if the RCO has special-enough cybersecurity needs to warrant its own cybersecurity program (distinct from their parent organization's cybersecurity program), then the RCO also needs a cybersecurity budget.

¹⁴¹ *See also* **Must 5 (Leadership)** for more details on the role leadership plays in major organizational decisions.

that cybersecurity is simply “being handled” by their IT (or some other department).

Moreover, establishing cybersecurity budgets provides a greater degree of transparency and rigor to the organization’s resource tracking. Explicitly calling out a central cybersecurity budget can help expose the hidden costs being placed on other organizational departments when they are left to deal with cybersecurity issues in a decentralized manner. This more realistic resource tracking can be used to support claims to stakeholders that the cybersecurity program has been given appropriate prioritization within the organization.

Finally, budgets reflect an organization’s commitment to cybersecurity and provide a valuable metric for evaluating and improving an organization’s cybersecurity program. Budgets can be compared between peer institutions or with other business sectors to gauge the organization’s cybersecurity program.¹⁴²

The Roadmap

This section describes the steps needed for an RCO to establish and maintain a cybersecurity budget. The process of developing a cybersecurity budget is fairly straightforward, but will vary greatly depending on the RCO’s usual budgetary processes. In short, the organization should employ its usual budgeting process and also do so for cybersecurity as a separate budget area. The following is generalized guidance that should be applicable regardless of budgetary process:

→ Start Early.

Budgeting for a cybersecurity program is most effective when it is started early in the organization’s lifecycle. Starting early in the lifecycle not only helps to make cybersecurity a part of the organization’s culture, but also demonstrates the priority of cybersecurity to other stakeholders. Moreover, starting an organization with a cybersecurity budget allows the organization to think strategically about its cybersecurity investments, focusing on long-term prevention and resilience rather than in a reactive manner to each new vulnerability or incident.

→ Consider Separating Cybersecurity From Your IT Budget.

Many organizations have a history of embedding cybersecurity within information technology departments. However, this is not a requirement and can lead to significant problems for some organizations. (*See **Must 7 (Cybersecurity Lead)** for a discussion of the factors impacting where cybersecurity should be located within the organization.*) Similarly, cybersecurity budgets need not be embedded within or dependent on the overall IT budget. Cybersecurity is an RCO-wide consideration, impacting the entire organization’s business processes, and its decision making should be independent of any one department.

¹⁴² Comparing budgets between institutions can be challenging if the organizations’ disagree about what costs should be included in the cybersecurity budget. For more details, *see* Common Challenge & Recommendations: “What Goes in a Cybersecurity Budget?” and “IT Budget vs. Cybersecurity Budget.”

→ Cybersecurity Lead Makes the First Proposal.

We recommend that the initial cybersecurity budget proposal in each fiscal cycle be made by the RCO's cybersecurity lead (*e.g.*, CISO), rather than a parent department or by the organization's overarching budgeting process. Having cybersecurity professionals make the first proposal sets the tone for the overall discussion, particularly when dramatic shifts in budgeting are necessary. Although RCOs may be tempted to have new budgets default to copies of previous years' budgets, with only basic adjustments up or down based on broader trends, this is not always an appropriate approach for cybersecurity. Cybersecurity is a dynamic and rapidly changing arena of strategic value to the organization, and cybersecurity programs require a considered approach to account for changes in the threat, regulatory, and technological landscape and to address actual or potential incident mitigation.

→ Benchmark Your Budget.

Cybersecurity benchmarking research on cybersecurity budgets commonly calculates budgets in a relative way, rather than in total dollars spent. The most commonly used benchmarking data compares either cybersecurity budget as a percentage of IT budget or cybersecurity budget as a percentage of the total organizational budget. Relative budgets provide a more consistent number across organizations of different sizes and in different business sectors. Although an RCO's baseline budget will invariably be in dollars (or other relevant currency), it is valuable to contextualize your organization's cybersecurity budget relative to some larger organizational budget.¹⁴³ Doing so allows for more effective benchmarking of cybersecurity spending relative to peer organizations and to broader guidance on cybersecurity spending generally.

Note, although cybersecurity benchmarking research often looks at cybersecurity budgets as a percentage of total IT budgets, this should not be interpreted as stating that cybersecurity is necessarily a part of IT. Rather, the total IT budget is a rough heuristic for establishing the similarities between organizations, and so this is used as a baseline to provide more useful comparisons between organizations of different sizes and sectors.

Common Challenges & Recommendations

This section describes some common **Must 12** challenges and offers recommendations on how to overcome them.

→ Budget Size.

One of the biggest challenges facing RCOs is determining what size budget is appropriate for their needs. Because organizational needs differ, there is no hard-and-fast calculation to use when provisioning resources for cybersecurity.¹⁴⁴ Benchmarking data consistently shows differing

¹⁴³ Russell, Jackson, Cowles; "Cybersecurity Budgeting: A Survey of Benchmarking Research and Recommendations to Organizations." 10 June 2016. <https://scholarworks.iu.edu/dspace/handle/2022/22289>.

¹⁴⁴ It is important to note that RCOs first establishing their cybersecurity program may see value in allocating a larger investment of resources, accounting for so-called "up front" costs (*i.e.*, one-time costs that are different from the *Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0*

cybersecurity budgets depending on organizational size and business sector, reflecting the differing needs of differently situated organizations.¹⁴⁵ Overall, benchmarking data shows security budgets vary quite substantially, ranging from 3% to 12% of IT budget.¹⁴⁶ Moreover, organizations that rely at least partially on a parent organization for IT and/or cybersecurity functions will have to take this into account when calculating their own budgets.¹⁴⁷ For a science-specific example, a 2016 NSF Cybersecurity Summit examination of public data on DOE open science laboratory cybersecurity spending¹⁴⁸ indicated that approximately 0.5% of the overall lab budget and 8-12% of the size of the IT budget (excluding scientific IT such as compute farms) were spent on cybersecurity.¹⁴⁹

For more discussion on the question of “how much” to budget for cybersecurity, *see* **Must 11 (Adequate Resources)**.

→ IT Budget vs. Cybersecurity Budget.

A recurring challenge when crafting cybersecurity budgets is distinguishing cybersecurity-specific practices from general IT practices. A number of cybersecurity practices are also just good IT practices, such as patching, configuration management, and code signing. Indeed, benchmarking data shows that organizations do not even agree on whether personnel labor should be included in cybersecurity budgets.¹⁵⁰ Since there is no clear answer to this question, differentiating IT and cybersecurity costs will likely depend on the specific culture of the organization. However, as the field of cybersecurity matures more broadly, organizations will likely begin to identify best practices with regard to organizational structure and cybersecurity budgeting, which will help organizations organize cybersecurity and IT programs more effectively.

→ What Goes in a Cybersecurity Budget?

Although **Must 12** is primarily concerned with the existence of a cybersecurity budget, it can be helpful to have some understanding of what types of activities are typically included in a cybersecurity budget. Although there is no established standard regarding what activities *must* be included in a cybersecurity budget, a basic breakdown would likely include cybersecurity

traditional year-to-year operational costs). However, these “up front” costs could also be spread over multiple budgetary cycles, if the RCO is unable to commit additional resources when first establishing their program.

¹⁴⁵ Smaller organizations consistently report higher cybersecurity budgets as a percentage of IT budget than larger organizations. This is likely due to the greater efficiencies that large organizations benefit from when investing in cybersecurity requirements, and due to minimum personnel requirements to cover the required areas of expertise.

¹⁴⁶ Russell, Jackson, Cowles; “Cybersecurity Budgeting: A Survey of Benchmarking Research and Recommendations to Organizations.” 10 June 2016. <https://scholarworks.iu.edu/dspace/handle/2022/22289>.

¹⁴⁷ Understanding how duties are shared between a parent organization and the RCO implicates a number of other **Musts**, most notably **Must 2 (Stakeholders & Obligations)** and **Must 15 (Baseline Control Set)**, (particularly Step 3 of the Roadmap).

¹⁴⁸ <https://static1.squarespace.com/static/5047a5a6e4b0dcecada15549/t/57b4b32dd2b857a1b6827a7f/1471460142220/Cybersecurity+Budgets+NSF+Summit+2016.pdf>.

¹⁴⁹ These numbers reflect the cost of being FISMA-compliant and are only the programmatic costs and do not include any operational costs such as continuous monitoring of the network, vulnerability scanning and patching, or the personnel to perform operational activities.

¹⁵⁰ *See, e.g.*, 2017 NSF Community Cybersecurity Benchmarking Survey Report, <https://scholarworks.iu.edu/dspace/bitstream/handle/2022/22171/2017%20Community%20Survey%20Report.pdf?sequence=2&isAllowed=y>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

For Public Distribution

management (*i.e.*, governance), operations, architecture, and incident response. The inclusion of personnel will likely vary depending on the size of the RCO's security program and on how many dedicated cybersecurity personnel are employed (as compared with IT personnel with cybersecurity responsibilities). Ideally, the budget would also cover adequate training for the cyber FTEs. Note, the discussion of how much to spend on individual cybersecurity activities is more fully discussed in **Must 11 (Adequate Resources)**.

Ultimately, however, the decision of what to include in the cybersecurity budget will be left to the RCO's leadership, taking into account their mission, culture, and established business practices.



Must 13: Personnel

Organizations must allocate personnel resources to cybersecurity.

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.

While all members of an organization have responsibility for cybersecurity at some level, a cybersecurity program must have personnel formally dedicated to developing, implementing, operating, and improving the program. This includes technical and non-technical roles that are resourced with organizational personnel to successfully support the mission. Cybersecurity personnel include people engaged in one or more of the following activities:¹⁵¹

1. Proactively protecting the organization from cyber threats and preventing the occurrence and recurrence of incidents.
2. Monitoring security operations and actively hunting for and detecting cybersecurity threats and vulnerabilities.
3. Responding to cybersecurity incidents and recovering to minimize impact.
4. Managing program governance, compliance, training and education, and risk communication management.

Personnel with cybersecurity responsibilities may be full-time cybersecurity employees, or may have a fractional full-time equivalent that includes both cybersecurity and non-cybersecurity responsibilities (*e.g.*, 0.5 Full Time Employee (FTE) allocated to cybersecurity and 0.5 FTE allocated to system administration). How each organization handles cybersecurity personnel resourcing can depend on its size, complexity, budget, and other factors.¹⁵²

Why is this a Must?

Cybersecurity is a constantly and rapidly evolving field, and tackling such a moving target requires dedicated organizational time and energy. This is challenging, if not impossible, without organizational personnel who are explicitly allocated with cybersecurity responsibilities. Formally allocating, measuring, and managing personnel effort is a key part of any successful cybersecurity program. It allows an organization to gain insight into the budgetary footprint of cybersecurity, which in turn makes it possible to apply the investment more strategically. It also gives cybersecurity leads more control over personnel with defined roles and responsibilities, allowing them to do their

¹⁵¹ Structuring the Chief Information Security Officer Organization (2015), https://resources.sei.cmu.edu/asset_files/TechnicalNote/2015_004_001_446198.pdf.

¹⁵² See **Must 11 (Adequate Resources)**.

job more effectively and thereby better protecting the organization. Furthermore, organizations are typically already performing cybersecurity activities, whether they are allocated or not. This includes personnel carrying out security functions for the organization, regardless of whether this is formally recognized in their role. Examples include IT staff managing incidents, non-IT staff patching software on a specialized instrument, and even researchers running antivirus scans on their desktops. Formally allocating these personnel resources to cybersecurity helps account for and consolidate this effort.

The Roadmap

This section describes the steps needed for a research cyberinfrastructure operator (RCO) to allocate personnel resources to cybersecurity. Allocating personnel effort to cybersecurity is not as straightforward as it may sound, especially for RCOs. This roadmap is designed to provide guidance and help RCOs ask the right questions to make the task easier. Step 1 begins with an exercise to determine the status quo of personnel allocation, followed by Step 2, which gauges the number of cybersecurity FTEs the organization needs. Step 3 then covers how to allocate cybersecurity effort. Finally, Step 4 talks about reevaluating effort to keep pace with organizational and environmental changes.

→ Step 1. Determine Current Cybersecurity Effort.

Before personnel resources are allocated, the organization should take a baseline of its current FTE effort on cybersecurity, both formal and informal. This includes creating an inventory of cybersecurity FTEs and locating them in the organization. For organizations with full-time cybersecurity personnel, the task will likely be easier to accomplish, but many RCOs do not fall in that category. A common challenge seen in RCOs is that IT staff are handling cybersecurity as an unacknowledged part of their job. Even in cases where IT personnel are explicitly tasked with cybersecurity, there may be complications such as:

- Cybersecurity being resourced or inherited from a parent organization. How this effort is counted is likely to be RCO specific and depend on factors such as the RCO financially contributing to the parent or local Human Resources practices, but it should be acknowledged formally by the cybersecurity program.
- Collaborations and virtual organizations where multiple organizations share the cybersecurity burden. Particularly interesting are cross-organization activities such as the RCO personnel securing a system located at a different organization.
- Use of cloud services. While the cloud provider might cover infrastructure security, the complexity of the cloud can make it difficult for the RCO to know where vendor effort ends and the RCO effort begins.
- Compliance. Many cybersecurity efforts have a compliance component that requires significant effort by non-security personnel, *e.g.*, by the general counsel and the compliance office.

Complexity notwithstanding, there *are* ways to tackle the inventory process. A good place to start is to work with the IT staff to learn about who is involved in which cybersecurity activities and in what

part of the organization.¹⁵³ This will help determine where the cybersecurity effort is currently going. It may also uncover other interesting tidbits such as IT already sharing certain security functions with the researchers, or, researchers spending time securing a system at another institution. This is important information to have in and of itself.

To document the existing effort, determine roughly what percentage of time people are spending on cybersecurity, and consider it as a partial FTE effort.¹⁵⁴ The results may not be precise, but even a rough estimate is better than nothing. A more accurate estimate may be possible by getting into the weeds and using system and application logs to determine time spent in cybersecurity activities such as patching, upgrading security software, and analyzing logs.

Another source of help may be a peer organization that has a cybersecurity program and has conducted an effort inventory exercise. Their approach may offer useful tips and guidance. Finally, for RCOs with a parent organization, the parent organization's CISO, CIO, or human resources offices may be able to provide insights into whether and how other areas within the organization are accounting for its cybersecurity effort.

→ Step 2. Determine the Effort Needed.

Once it is known where the current cybersecurity effort is being expended, the next step is to estimate personnel resources it will take to run an effective cybersecurity program for the organization. This is likely to be the most unwieldy step since it depends on many variables, some of which may be unknown. Trying to find guidance on how to resource personnel for a cybersecurity program is also riddled with issues:

- There is limited research in the area.
- From the few studies and surveys that exist, the overwhelming answer seems to be “it depends.”^{155,156} The primary dependencies are the type of organization, size, and complexity.
- The studies use data from the corporate sector and large organizations. It is unclear to what degree these results are applicable to RCOs.
- RCOs that have access to IT consulting firms (*e.g.*, Gartner)¹⁵⁷ can leverage this resource to provide professional feedback on the analysis of effort needed and effort allocation.

Despite these shortcomings, the studies provide the best evidence available and benchmarks that are simple to use. They cast the requisite cybersecurity effort in convenient, easy to understand, and quantitative terms¹⁵⁸ (*see also* **Must 11 (Adequate Resources)** and **Must 12 (Budget)**). They can be summarized as follows.

¹⁵³ The NICE Cybersecurity Workforce Framework provides a comprehensive listing of work roles organized by functional categories and specialty areas. This is a good resource to identify cybersecurity activities associated with work roles. <https://niccs.cisa.gov/about-niccs/nice-cybersecurity-workforce-framework-work-roles>.

¹⁵⁴ Is there a resource planning formula?, <https://meisterplan.com/blog/is-there-a-formula-for-resource-planning/>.

¹⁵⁵ <https://www.nuharborsecurity.com/information-security-staffing-guide>.

¹⁵⁶ Structuring the Chief Information Security Officer Organization (2015), https://resources.sei.cmu.edu/asset_files/TechnicalNote/2015_004_001_446198.pdf.

¹⁵⁷ <https://www.gartner.com/en/about>.

¹⁵⁸ Structuring the Chief Information Security Officer Organization (2015), pp. 16-17, https://resources.sei.cmu.edu/asset_files/TechnicalNote/2015_004_001_446198.pdf.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

- Security FTEs as a percentage of IT FTEs: 5.2% (This percentage excludes staff responsible for business continuity and disaster recovery.)
- Approximately 1 FTE is dedicated to cybersecurity per 5000 networked devices (workstations, switches, firewalls, servers, etc.)
- Between 3% and 11% of the total IT budget is dedicated to cybersecurity

Finally, using the information in Step 1, it may be possible to reduce the FTE count by considering inefficiencies or exploring alternative ways of performing the same tasks by shifting resources (*i.e.*, purchasing a tool that makes processes less time-consuming).

→ Step 3. Allocate Effort.

Armed with the results of Steps 1 and 2, it should now be possible to identify the gap between the effort needed and the effort available.¹⁵⁹ With this estimate comes the task of determining how to best allocate this effort. Here are some key questions to consider to help guide the exercise.

- **How should the effort be distributed?** The minimum necessary requirement for a cybersecurity program is to have a cybersecurity lead who has been assigned responsibilities as per **Must 7 (Cybersecurity Lead)**. Other responsibilities may be best distributed among multiple people. To get an idea of the range of possible cybersecurity roles and responsibilities, consider the NICE Cybersecurity Workforce Framework¹⁶⁰ as a reference. For example:
 - System Security Analyst - Responsible for the analysis and development of the integration, testing, operations, and maintenance of systems security.
 - Security Architect - Designs enterprise and systems security throughout the development lifecycle; translates technology and environmental conditions (*e.g.*, law and regulation) into security designs and processes.
 - Cyber Defense Analyst - Uses data collected from a variety of cyber defense tools (*e.g.*, IDS alerts, firewalls, network traffic logs) to analyze events that occur within their environments for the purposes of mitigating threats.
 - Security Control Assessor - Conducts independent comprehensive assessments of the management, operational, and technical security controls and control enhancements employed within or inherited by an information technology (IT) system to determine the overall effectiveness of the controls.
 - Forensics Analyst - Conducts deep-dive investigations on computer-based crimes establishing documentary or physical evidence, to include digital media and logs associated with cyber intrusion incidents.
 - Cyber Legal Advisor - Provides legal advice and recommendations on relevant topics related to cyber law.
 - Cyber Defense Incident Responder - Investigates, analyzes, and responds to cyber incidents within the network environment or enclave.

¹⁵⁹ We only mention the gap here because an excess of cybersecurity FTEs is inconceivable to us, at least in research organizations.

¹⁶⁰ <https://niccs.us-cert.gov/nice-cybersecurity-workforce-framework-work-roles>.

- What does it mean to allocate fractional FTEs to cybersecurity? Each cybersecurity duty or role requires a certain amount of effort. Fractional effort allocation means assigning discrete cybersecurity duties (*e.g.*, incident response) or roles (*e.g.*, security engineer) to a person for only a fraction of their total FTE time. The cybersecurity lead role may itself be a fractional FTE if the lead has other duties, for example, system administration.
- When does fractional allocation make sense? The first scenario is when there are simply not enough personnel resources to carve out a full FTE for cybersecurity. Fractional FTEs may also be useful when cybersecurity skills are distributed across the organization and best left there. (Both scenarios are likely to be true for many RCOs.) Since there is already a cybersecurity effort whether or not it is formal, you should have a rough idea of what it is and how it is distributed from Step 1. Start by formally acknowledging and taking this effort into account.
- What fraction of the cybersecurity lead should be dedicated to cybersecurity? In an ideal world, each organization would have a full-time, dedicated cybersecurity lead, but this may not always be feasible or necessary. Determining optimal effort allocation for the lead runs into the same issue as Step 2, namely, it depends on organization type, size, and complexity. If the cybersecurity lead is the only dedicated cybersecurity person, the benchmark in Step 2 could be used to calculate rough fractional effort, *i.e.*, 5-10% of the IT FTEs. Another possibility is to start with this number while having the cybersecurity lead document actual time spent on cybersecurity, and adjust.
- What if the cybersecurity lead is the only dedicated cybersecurity person? This is the case for some RCOs, and the priority should be for the lead to carefully optimize their own FTE effort.
- Who should be tasked with allocating effort? Ideally, the cybersecurity lead should allocate personnel effort within the program, as the cybersecurity program manager per **Must 7 (Cybersecurity Lead)**. However, this may not be possible if the lead does not have the proper authority or if the organization has a different personnel allocation model, for instance, human resources making such decisions. RCO leadership should ultimately decide the course of action best suited to the mission's needs.
- How should effort allocation be coordinated? RCO leadership may empower the cybersecurity lead to allocate effort, but coordinating and putting it into practice requires skill. One thing is certain—it is unlikely to be achieved by fiat, especially for RCOs. If cybersecurity effort is distributed across multiple groups, diplomacy and cooperation will be needed to effect a positive outcome. The lead should communicate with group leaders and negotiate who within their group can be part of the organizational cybersecurity program without significantly impacting research productivity.

Finally, RCOs may consider establishing a cybersecurity policy specifically addressing personnel resource allocation. For example, specify the required effort allocation for each role outlined in the roles and responsibility section.

→ Step 4. Reevaluate.

Reevaluation is a key piece of a cybersecurity program, and personnel resource allocation is no exception. We recommend a regular exercise, at least annually, to evaluate needs going forward, asking “What has changed since the last evaluation?” Use the trigger events outlined in **Appendix B** as a guide for reevaluations.

Common Challenges & Recommendations

This section describes common **Must 13** challenges and offers recommendations on how to address them.

→ Tackling Multi-institutional Projects.

RCOs often support collaborative research involving multiple, geographically dispersed institutions, with varying degrees of cybersecurity capability. One approach that has been tried successfully in a number of virtual research organizations is to assign a project-specific cybersecurity lead to oversee cybersecurity and to manage the distribution of responsibility. The cybersecurity lead can be tasked with assessing where personnel gaps exist, determining the effort required to fill it, and developing strategies to distribute resources among the participating institutions.

→ Requesting Personnel Resources.

In an ideal world, not only would an RCO have senior leadership involved as per **Must 5 (Leadership)** and a designated cybersecurity lead as per **Must 7 (Cybersecurity Lead)**, it should also be able to submit a formal request to fill the gap, with a reasonable likelihood of getting most or all of what it needs. In the real world however, leadership involvement and a cybersecurity lead may be within the realm of possibilities, but getting additional FTEs is exceedingly difficult, if not altogether impossible. To have a fair shot at success, there needs to be a solid, written justification. It may fail, but at least there will be documented evidence of the formal request, a useful artifact in case there is a breach and an ensuing investigation later. Here are some possible arguments to use, if applicable:

- **The Bottom Line:** Determine the absolute minimum effort necessary to ensure mission continuity and show that it exceeds what you have at present.
- **Growth:** Determine the effort necessary to handle known, imminent growth, for example the \$10 million contract the organization just received, and show that it exceeds what you have.
- **Compliance:** Determine the effort necessary to comply with rules and regulations increasingly included in sponsor terms and conditions, especially in that \$10 million contract, and show that it exceeds what you have.

→ Picking the Right Personnel.

While ancillary to the core issue this Must addresses (dedicating personnel resources to cybersecurity), optimally allocating cybersecurity effort also requires knowing who to allocate this effort to. Ultimately, a cybersecurity program is only as good as its people.

Cybersecurity requires a wide variety of skills.¹⁶¹ Many of its core activities are technical, so it needs at least some personnel who are technically competent (or trainable). Ideally, they would have a deep understanding of network protocols and devices, operating systems and web servers, application and database security, and (potentially) secure configuration of cloud-based services. In reality, RCOs often lack deep cybersecurity expertise due to their research-focused mission, and acquiring it through new hires is difficult due to budgetary and market pressures. Even if the personnel have the requisite technical expertise, it may not be enough, since cybersecurity also requires soft skills. This includes communication skills, teaching skills, and negotiating skills.¹⁶² Communication skills are required to make technical and non-technical staff aware of how their action or inaction affects the organizational mission, and to communicate the organizational impact of risk decisions to information asset owners and senior management.

The most effective cybersecurity staff are those who can form good personal trust relationships with other professionals in various "trust" groups and collaborative bodies. These professional contacts are an essential part of gaining up to date intelligence on threats and learning how best to deal with them. Negotiation skills are also important since they are needed to get researcher buy-in for security initiatives. Finally, the right personnel should be reliable, trustworthy, resilient, and willing.

We recommend using organizational personnel where possible. Making major programmatic and risk acceptance decisions requires deep organizational knowledge. Insiders are likely to be much more invested in the organization than outsiders, such as a vendor. Another, important reason for picking internally is to cultivate and retain cybersecurity expertise in-house. Cybersecurity internships may be another mechanism for personnel resourcing which may pan out into permanent personnel.

→ Retention.

Retaining cybersecurity professionals is as much of a challenge as hiring. Many professionals move from job to job for higher pay and promotions while others seek a better work culture/environment.¹⁶³ Many research institutions may not be able to compete with the private sector when it comes to monetary compensation. RCOs should consider building a retention strategy that does not rely solely on salary or compensation.¹⁶⁴ At the onset, promote the organization's culture and mission of

¹⁶¹ The Future Cybersecurity Workforce: Going Beyond Technical Skills for Successful Cyber Performance, <https://www.frontiersin.org/articles/10.3389/fpsyg.2018.00744/full>. This article discussed the characteristics of a successful cybersecurity workforce.

¹⁶² An ISC2 study identified the following as some of the most important qualifications, in decreasing order of priority: relevant experience, advanced cybersecurity knowledge, certifications, work experience, strong soft skills, knowledge of regulations, and cybersecurity related undergraduate degree, <https://www.isc2.org/-/media/ISC2/Research/2018-ISC2-Cybersecurity-Workforce-Study>.

¹⁶³ <https://www.infosecurity-magazine.com/blogs/talent-retention-emphasis-1-1/>.

¹⁶⁴ <https://www.xaasjournal.com/3-tips-for-finding-and-retaining-cybersecurity-talent/>.

supporting science as a benefit. Highlight the organization's total compensation and benefits package, (*e.g.*, healthcare, retirement, PTO), which may be more desirable than those offered by private sector employers. Also, highlight other non-monetary benefits such as alternate work schedules, reduced or free education and training, and flexibility to work on different projects. Use the guidance provided in the previous recommendation and avoid over-taxing personnel with excessive responsibilities, long hours, and limited time to rest.¹⁶⁵

Organizations should invest in measures that enable retention and job satisfaction. Setting aside some funds for professional development is one strategy that can promote both. Encouraging and facilitating activities that promote learning, such as a mentorship program, security training, and opportunities for networking, (*e.g.*, attending conferences), are crucial to fostering a progressive and successful cybersecurity program. Being sensitive to the pressures a dynamic and fast-changing environment can impose on cybersecurity personnel and presenting them as potential opportunities for learning can also help increase job satisfaction and allow them to gain highly marketable skills. For instance, tabletop security exercises and red and blue teaming are both fun and useful exercises. Also point out advantages of working in research and higher education such as campus life, relaxed atmosphere, great benefits, and serving a role in advancing the state of knowledge in job ads. Finally, providing avenues for career advancement within the organization is always a strong motivation for personnel to stay with the RCO.

¹⁶⁵ *Ibid.*



Must 14: External Resources

Organizations must identify external cybersecurity resources to support the cybersecurity program.

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.

Musts 12 (Budget) and **13 (Personnel)** emphasize minimum necessary attention to internal resources to support the adequacy requirement in **Must 11 (Adequate Resources)**. **Must 14** emphasizes that internal resources alone are insufficient to support a competent cybersecurity program. Organizations should not “go it alone” when developing a cybersecurity program. External cybersecurity resources help prevent “reinventing the wheel” and help the organization to utilize available resources more efficiently. External cybersecurity resources include parent¹⁶⁶ and peer organizations, consortia that provide services, security consultants, and commercial vendors as they can provide both general programmatic recommendations or specific services and expertise, as needed. Resources also include open source and commercial^{167,168} tools (*e.g.*, for penetration testing, malware detection, and log monitoring and analysis). It is important however to remember that effective use of these external resources often involves substantial personnel costs both initially (*e.g.*, to migrate to a cloud environment) and on an ongoing basis (*e.g.*, monitoring using an intrusion prevention tool).

External resources are essential to the success of a cybersecurity program, but only when used wisely. The key question to ask before leveraging them is, “Will this make my job easier so I can focus on other priorities?”

¹⁶⁶ Many RCOs have parent organizations and these parent organizations may be the primary go-to “external” resource. There is considerable variety in those relationships in terms of operational independence, similarity of mission needs, communication, common culture, and legal relationship. As such, the distinction between an “internal” and “external” resource may be fuzzy. This Framework Implementation Guide is predicated on the idea that an RCO, regardless of the existence of a parent organization, may benefit from having its own cybersecurity program even if that program *heavily* leverages policy, services, and other resources that sit at the parent organization level. Even in these circumstances, the reasons for a defined RCO cybersecurity program addressing all **Musts** can include the RCO’s distinctive missions, stakeholders, and requirements, novel technological resources, and different threat profiles.

¹⁶⁷ “[W]here unique skillsets are needed, an outside firm may be engaged to assist with the cyber risk assessment, implementation of resilience measures, and/or periodic assessments of the effectiveness of the program. ... [I]f an organization experiences a significant cybersecurity incident or breach, outside expert assistance may be needed. ...” <https://www.coso.org/Documents/COSO-Deloitte-Managing-Cyber-Risk-in-a-Digital-Age.pdf>, pg. 7.

¹⁶⁸ The RCO community typically favors freely-available, open source solutions, but they are not necessarily the optimal choice in each case.

Why is this a Must?

Due to the complexity and uncertainties involved, organizations in general and research cyberinfrastructure operators (RCOs) in particular do not have sufficient expertise, personnel time, or funding to develop and maintain an effective cybersecurity program entirely in house. Organizations which have specialized in cybersecurity services often have well-developed specialist knowledge and expertise in the domain they serve and are capable of providing services of a higher quality than would be possible for an RCO to achieve by developing an in-house solution. In addition, making use of applicable external resources allows an RCO to avoid duplication of effort. By leveraging existing solutions offered by third-parties, an RCO can avoid retreading the same path where others have already done the work. Finally, external resources allow the RCO to focus time and limited resources on those parts of the organization's cybersecurity mission which require the unique attention or expertise of personnel with detailed knowledge of the organization, and on those problems which are more unique for the RCO for which there are not already well-understood solutions.

The Roadmap

This section describes the steps needed for an RCO to identify external cybersecurity resources to support the cybersecurity program. There is a wide variety of external cybersecurity resources. To make it easier to leverage them, we recommend the following, stepwise, iterative process.

→ Step 1. Survey the Landscape.

The first step is to know what is available. This awareness plays a critical role while planning for security from the earliest point possible to effectively balance the costs and benefits associated with these resources, as well as specific product/vendor choices, especially in light of the existing cybersecurity skills shortage.

A natural place to start is the RCO's own or parent organization's resources. It is often possible to leverage existing enterprise policies and procedures (*e.g.*, incident response) and organizational resources such as the following:

- Assistance with assessments
- Security training
- Security services (*e.g.*, monitoring, firewalling)
- Forensic services
- Internal audit
- Help with relevant laws and regulations (*e.g.*, DFARS, HIPAA, state data breach, and data retention statutes)

Moving beyond the local, there are a wide array of external resources available to an RCO, which can be both a boon and a bane. Having choices is great, but the cybersecurity marketplace is so huge that an RCO can get lost in it. To help get started, we have broadly categorized a number of options RCOs should be aware of.

1. Research and education (R&E) peers. Whether you are an RCO for a research project, a university, or a research organization/facility, peers with a profile similar to yours can be an invaluable resource for information and resources. Use opportunities such as working groups, meetings, and conferences¹⁶⁹ to make connections.
2. Consortia.¹⁷⁰ There is an assortment of these operating in the R&E space that provide cybersecurity resources and services such as federated identity, security event monitoring and alerting, threat intelligence, and consulting.¹⁷¹
3. Cybersecurity associations and trade groups. There are a variety of such organizations¹⁷² providing information and other resources with and without membership fees that cover areas such as vulnerabilities, threat intelligence, and cloud security.
4. Open source software. Thanks to both small and large open source projects, there are many, freely available cybersecurity tools,¹⁷³ often with a large, supportive user community behind them.
5. Commercial vendors. The marketplace for cybersecurity services and tools is large. Vendors can provide software, managed services, consulting, monitoring, incident response, penetration testing, and much more. Since the price tags of these offerings can be substantial, we recommend researching reviews online and contacting peers to gather information about appropriate and affordable third party options. Another, important strategy is to leverage resources that come with substantial, built-in cybersecurity and compliance,¹⁷⁴ for instance cloud vendors. An important caveat to be mindful of when considering a commercial vendor or tool is the time and effort it takes for procurement, especially if speed is of essence. While it may first appear that leveraging a commercial tool will accelerate implementation, leadership approvals, vendor agreements, vendor security assessments, etc. can take substantial time. Sometimes a vendor is not even willing to engage in a legal dialog.

→ Step 2. Assess Opportunities and Needs.

Armed with the knowledge of available options, the next step is to identify potential candidates by considering local needs, environment, and realities. Depending on the organization, this can be straightforward, for instance picking an easy to use, open source tool for a small RCO with a limited budget, or challenging, for example when there is a bewildering array of vendors providing a similar product or service.

¹⁶⁹ For instance, the NSF Cybersecurity Summit (<https://www.trustedci.org/2020-nsf-summit>) and EDUCAUSE Security Professionals Conference (<https://events.educause.edu/security-professionals-conference/2020>).

¹⁷⁰ Entities that have been formed by organizations pooling their resources together to receive services (typically).

¹⁷¹ Some examples include the ResearchSOC and OmniSOC (security monitoring, alerting), REN-ISAC (peer security assessments, threat intelligence), InCommon (federated identity), and Trusted CI (security assessments, consulting).

¹⁷² For example CERT, ISACA, OWASP, Cloud Security Alliance, and CIS.

¹⁷³ For example Zeek, ELK, Snort, ClamAV, OSSEC, etc.

¹⁷⁴ For example AWS GovCloud or Azure Government Cloud Computing (GCC).

To help narrow down the choices, we recommend taking into account factors such as the following:

1. Types of resources that fit existing gaps or add benefits to the program that simply cannot be fulfilled by the RCO itself (*e.g.*, threat sharing organizations that have capabilities beyond the reach of any single organization.)
2. The nature of the organization's relationship with the resource (*e.g.*, paid/unpaid; direct or at-a-distance; membership model).
3. The nature of the costs of leveraging it (*e.g.*, money, time, attack surface expansion).
4. Location (*e.g.*, on-prem vs. cloud).
5. The value of the relationship in the long term.
6. Compliance (*e.g.*, vendors that offer FedRAMP¹⁷⁵ compliant cloud services).
7. Vendor reputation.

→ Step 3. Pick Judiciously.

Leveraging external resources always comes at a cost, which may range from nominal, *e.g.*, subscribing to and reading security alerts, to substantial, *e.g.*, implementing an Intrusion Detection System (IDS) that requires major FTE effort. Organizations typically, if not universally, do not have adequate internal resources (informational/intelligence, human, technological) to run a competent cybersecurity program. The cybersecurity problem is simply too complex for any organization to tackle in a vacuum. Consequently, some advanced thinking about trade offs goes a long way toward optimizing the use of external resources and avoiding problems later. We recommend that you ask yourself the following questions prior to committing to a resource:

1. Do I really need it?
2. Can an existing, in-house solution be used or enhanced to do the same thing?
3. Do I have the necessary resources to implement and manage it?
4. How will it affect the users?
5. What is the cost vs. benefit?
6. Is it sustainable in the long run?

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.

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.

¹⁷⁵ <https://www.fedramp.gov/program-basics/>.

→ Step 4. Engage & Iterate.

Once a decision is made to leverage an external resource, the next step is to acquire and implement it. This seems simple enough but still requires some planning to execute properly. Potential items to consider include the following:

1. Assigning personnel
2. Implementation schedule
3. Transition to a new resource
4. Staff and user training

Using an external resource, especially one that is new, provides an opportunity for review: how well the resource met expectations, whether it was cost effective, how difficult it was to implement, user uptake, and so forth. The final step of **Must 14** then is to use this information to ascertain what could be done better the next time and to share this information with the community if and when the opportunity presents itself.

Common Challenges & Recommendations

This section describes common **Must 14** challenges and offers recommendations on how to address them.

→ Cost of External Services and Tools.

Sophisticated technology products not only have upfront costs but also require skilled expertise to install, configure, maintain, and monitor which necessitate increased local staff or external resources. Also, the upfront cost and ongoing expense for each solution may differ greatly. In order to help economize, we recommend the following strategies:

1. Coordinate with your own or parent organization (if mission requirements are compatible).
2. Leverage R&E peers and consortia.
3. Explore discounted pricing with the vendor (many offer them to research and non-profit organizations).

In addition to acquiring and deploying tools and services, outsourcing them could also be an option. While it may appear beyond the reach of RCO budgets, outsourcing can in certain cases be a wise (and less costly) choice, for instance when the cost of additional in-house personnel exceeds the outsourcing cost.

→ Compliance.

While in the past most types of research have been relatively untouched by compliance burdens such as those imposed by federal regulations, this is changing quickly. Grants, contracts, and data use agreements now routinely come with cybersecurity strings attached. Compliance is an especially difficult problem for research organizations owing to a lack of expertise. The good news is that the growing compliance burden is quickly adding to the number of RCOs and other R&E entities with

experience and expertise in the subject to rules, regulations, and use agreements. Most of these organizations are willing to share information and provide peer guidance,¹⁷⁶ making them an ideal first place to go, especially for RCOs facing compliance for the first time. Commercial vendors are also an option, albeit much more costly.

A different pitfall is a vendor self-professing compliance with regulations. RCOs without prior experience can be lulled into thinking the vendor to be a complete cure for compliance. Unfortunately, for most compliance regimes affecting research, vendor compliance will only be a piece of the puzzle and require plenty of due diligence from the RCO.

→ Vaporware, False Promises, and Hidden Costs.

Vendors often take advantage of customer's fear, uncertainty, and doubt (FUD) by inflating what their products can do or advertise a "sell it and forget it" solution. It is often challenging for the unsuspecting RCO to discern the bad from the good. Online research is a good place to begin to ensure the product delivers what it promises. Talking with peers and others with experience with the product is another avenue. Marketing speak presents another pitfall as the vendor may hide the in-house cost of deployment and maintenance. Quality vendors are always transparent with their products and support offerings.

→ Vulnerabilities from Outside Entities.

While vulnerabilities in security tools are well known,¹⁷⁷ any trust relationship with an external entity is not without increased danger. "In many cases, trust relationships with external organizations, while generating greater productivity and cost efficiencies, can also bring greater risk to organizations. This is addressed by strategies established by organizations that take into account the strategic goals and objectives of organizations."¹⁷⁸ In particular, ensure written agreements with outside entities, including users, provide for notification of compromises and/or vulnerabilities and security updates.

→ Evolving Landscape.

The threats, tools, and services are in a continual state of flux. RCOs need to monitor the environment if they are going to optimize use of the external resources available. "Organizations must operationalize governance processes to capture and evaluate potential changes that may alter their cyber risk profile. This includes, at a minimum, capturing prospective new and changing products and services, information technology and evolving digital strategies, business processes, mergers, acquisitions, and reorganizations, and laws and regulations."¹⁷⁹ Having cybersecurity personnel participate in conferences, cybersecurity email lists, and webinars aid greatly in monitoring the changing cybersecurity environment.

¹⁷⁶ For example, Trusted CI.

¹⁷⁷ <https://www.rack911labs.com/research/exploiting-almost-every-antivirus-software/>.

¹⁷⁸ NIST SP 800-39, pg. 25.

¹⁷⁹ <https://www.ciso.org/Documents/COSO-Deloitte-Managing-Cyber-Risk-in-a-Digital-Age.pdf> pg. 14.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

For Public Distribution

→ Managing Outsourced Cybersecurity Tasks.

Engaging third parties to provide cybersecurity services can be useful for RCOs with limited IT or cybersecurity resources. Third party services can be particularly important given the shortage of trained, experienced cybersecurity personnel available, but they do increase the workload in terms of managing the outsourced service(s). “However, in the event that tasks related to cybersecurity are outsourced, it is essential for the organization to perform the following:

- Maintain regular communication with the service provider for awareness of incidents
- Discuss new and potential threats as the organization’s business environment changes and cyber threat landscape continues to evolve
- Provide open communication lines for immediate escalation when a significant incident or breach occurs.”¹⁸⁰

→ Risk escalation.

A concept from project risk management -- risk escalation -- is “notifying a higher level authority. If a risk is outside of the scope of the project or the proposed response is outside the project manager’s authority, it would make sense to escalate the risk to a higher-level manager within the organization.”¹⁸¹ This concept may be helpful for RCOs.

Translated for use in this context, consider: **If a risk is outside the scope of the RCO’s mission or ability to act, or the proposed RCO response is outside the RCO’s authority, it would make sense to escalate the risk to a higher authority.** Specifically for RCOs, the risk escalation strategy may be built into the relationship between the RCO and its parent organization and/or funding agency(ies), where a primary strategy is to notify of the need for help in addressing risks. It presents an alternative to risk acceptance.

This strategy is relevant to cybersecurity in general because malicious cyber attacks are unlawful, may be outside the reasonable scope of a capability to defend against (*e.g.*, against advanced nation-state attackers), and organizations lack the lawful capacity to counterattack effectively. These attacks are the cases where escalation to law enforcement and national security functions should at least be considered.

This strategy presents a challenge when the higher authority is unable or unwilling to assist.

¹⁸⁰ *Ibid*, pg. 14.

¹⁸¹ Schwalbe, K. (2019). Information Technology Project Management. 9th ed.
Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0
For Public Distribution



Controls



Must 15: Baseline Control Set

Organizations must adopt and use a baseline control set.

Controls are specific administrative, technical, and physical safeguards and countermeasures applied to reduce cybersecurity risk. A baseline control set is a predetermined set of controls used as a default when selecting security controls for information assets. The baseline control set does not determine what security controls an organization must implement; rather, it provides a foundation from which an organization tailors control selection based on the needs of its mission. Baseline control sets vary in the number, specificity, and goals of the controls it describes. Baseline control sets may be legally imposed when handling specific types of data. In other cases, organizations can select a well-maintained control set that is based on evidence of what works to reduce cybersecurity risk.

Research cyberinfrastructure operators (RCOs) may have some discretion when determining which baseline control set or sets to adopt and use. However, with so many options and opinions, baseline control set selection can be challenging and contentious. Organizations are often torn between selecting the most appealing control set for the present and potentially future-proofing against future regulations. Baseline control set selection, therefore, benefits from eliciting help from cybersecurity experts, peers, and other available organizational resources. Often host or peer IT organizations are available to discuss this topic. Trusted CI is also available to aid with control set decisions. The Center for Information Security (CIS) Controls¹⁸² are widely regarded as authoritative and reasonable.¹⁸³ The CIS controls are 1) highly prioritized; 2) updated frequently; 3) described in sufficient detail for practitioners to implement them; and 4) developed by a collaborative and open process informed by a diverse group of cybersecurity practitioners.

Trusted CI recommends adoption of the CIS Controls¹⁸⁴ unless another baseline control set is legally or contractually required for selected information asset categories.

Baseline control sets aid organizational cybersecurity by setting out established best practices that organizations can rely upon when selecting controls. Although cybersecurity is a quickly evolving

¹⁸² <https://www.cisecurity.org/controls/>; for more detail including history, *see*

<https://www.sans.org/critical-security-controls>. We recommend using the latest version, and remaining apprised of updates. Direct link to CIS Controls Version 7.1:

https://drive.google.com/file/d/1-ge3sa_qmL-qCnYosMZXR2dMGGUgFymx/view. CIS Controls Version 7.1 is licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License (the link can be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

¹⁸³ The 2016 California Data Breach Report,

<https://oag.ca.gov/sites/all/files/agweb/pdfs/dbr/2016-data-breach-report.pdf>, states that failure to implement the controls described by CIS that apply to an organization's environment constitutes a lack of reasonable security. Further, it also states that multifactor authentication should be available for online accounts with access to sensitive information and that such information should be encrypted on portable devices.

¹⁸⁴ <https://www.cisecurity.org/controls/>.

field and organizations are likely to face unique or challenging risk management decisions, there is a great deal of solid advice available, particularly around commodity IT and general consensus best practices. These best practices are embodied in control sets, such as the CIS Controls. Effective, prioritized, and evidence-based control sets are the best economic choice in terms of the protection offered for the resources invested. The publication “Economics of Cybersecurity: A Practical Framework for Cybersecurity Investment”¹⁸⁵ identifies three key principles for making cybersecurity investment decisions: 1) Employ a baseline control set to protect against the majority of attacks that are generally of low to moderate sophistication; 2) Employ advanced security mechanisms to protect organizations’ critical functions and data against sophisticated threats; 3) Accept the risk of not protecting less mission-critical functions when cost of implementing security mechanisms is much higher.

Why is this a Must?

Adopting and using a baseline control set allows an organization to rally around a common language and structure for a set of controls designed to address common assets, threats, and effective defenses. It relieves the resource burden of *ad hoc* control selection and mitigates the major risks of *ad hoc* control selection: missing important, doable controls and wasting effort “reinventing the wheel.”

The Roadmap

This section describes the steps needed for an RCO to adopt and get started using a baseline control set or sets. Step 1 is about knowing yourself and knowing the available options: There are a lot of baseline control sets out there, and not all are created equal. Step 2 is the formal selection and adoption of the set(s). Step 3 is determining the relevance of the controls in a set, their current state of implementation, and making decisions regarding what actions are to be taken as time, resources, and priorities permit.

An RCO’s leadership (**Must 5**), cybersecurity risk acceptors (**Must 6**), and cybersecurity lead (**Must 7**) all have roles to play in aligning to **Must 15**.

→ **Step 1. Know what you need.**

Step 1 is to understand your organization’s mission and how the various options for baseline control sets intersect with your organizational needs. There are a number of baseline control sets available for organizations to select from, and these control sets often serve different purposes or different audiences, so the selection of a baseline control set requires both understanding what that control set was designed for and whether it is appropriate for your organization. Research cyberinfrastructure operators typically support a research mission that is highly collaborative, frequently international, and possesses unique security needs from what exists in other communities. That said, even if not required to do so, RCOs should not ignore the value of aligning to existing

¹⁸⁵ The Armed Forces Communications and Electronics Association (AFCEA)’s Cyber Committee has produced a useful and relevant publication: “The Economics of Cybersecurity: A Practical Framework for Cybersecurity Investment,” <https://www.afcea.org/committees/cyber/documents/cybereconfinal.pdf>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

control sets since they are likely already familiar to key stakeholders, and it simplifies the process of electing a baseline.

This selection process therefore has important interactions with **Must 1 (Mission Focus)**, **Must 2 (Stakeholders & Obligations)**, **Must 3 (Information Assets)**, and **Must 4 (Asset Classification)**. **Must 1** is important because it is where the organization identifies the role that cybersecurity plays in supporting the organization’s mission. **Must 3** is important because it is where the organization identifies the assets to which the control set will apply and **Must 4** guides the application of specific controls to categories of information assets. **Must 2** is important because an RCO’s decision may be influenced or constrained by laws, regulations, contracts, and parent organization policy, as well as by less formal preferences of stakeholders. For instance, control sets may be legally required for particular types of information or information systems (e.g., Controlled Unclassified Information (CUI) with NIST SP 800-171).

In most cases, however, organizations will not have a sweeping mandate, and will therefore need to understand the range of options. There are a number of control sets in the wild, ranging from ones envisioned as appropriate for a broad range of organizations to those tailored for particular communities, types of information, or technologies. These vary in specificity, number of controls, and how they are produced and maintained. We characterize a number of these control sets in Appendix C. Many controls are common to most baseline control sets, and control set “crosswalks” are available.¹⁸⁶

→ Step 2. Adopt the control set (or sets).

Once an RCO has sufficiently reviewed the options available to them, consistent with the RCO’s mission (**Must 1**), and stakeholder requirements (**Must 2**), the RCO has to formally select and adopt their baseline control set(s). The process of adoption involves the RCO formalizing its use of the baseline control set as its default when selecting controls. Final approval for the adoption of a baseline control set should rest with senior management. Adoption of the baseline control set should be formalized in policy: A master information security policy is a natural place to formalize control set adoption.¹⁸⁷

Trusted CI recommends the CIS Controls unless another baseline control set is legally or contractually required for selected information asset categories. The CIS Controls are updated regularly based on input from a diverse community of experts. The varied viewpoints of these diverse experts ensure the CIS Controls address current threats, are complete, and are prioritized by organizational type.

Note, however, that RCOs may find that they must or should adopt and use more than one baseline control set. For instance, an RCO might adopt the CIS Controls as their general, default set, but adopt or add on controls laid out in the Payment Card Industry-Data Security Standard¹⁸⁸ only for

¹⁸⁶ <https://www.auditscripts.com/download/2742/> will download a spreadsheet that is very comprehensive in terms of both the CIS controls and the frameworks covered.

¹⁸⁷ Check out Trusted CI’s **Master Information Security Policy & Procedures Template**, available at <https://trustedci.org/framework>.

¹⁸⁸ https://www.pcisecuritystandards.org/document_library?category=pcidss&document=pci_dss.


systems handling credit card information.

→ Step 3. Baseline the set.

Once an organization has adopted a baseline control set, we shift to what it means to “use” it. The basic process entails determining (a) the relevance of the individual controls based on the requirements of the categories defined in **Must 4**, (b) the current state of implementation, and (c) whether the current state is acceptable or warrants further action.

- (a) **Determine relevance.** A particular control may or may not be relevant for the organization, and if relevant, it may only be relevant for a particular type of asset. The bar for a control being relevant should be kept low. If a determination of “not relevant” is made, we encourage documenting the reason why. If the relevance is highly specific to particular types of assets or part of the RCO’s operation, we encourage documenting that detail.
- (b) **Determine current status.** Because control statements are often not prescriptive, the RCO should document its interpretation and how to address at a high level the control’s objective. Particularly for organizations in active operations, work may be required to determine the current state of a particular control’s implementation. The good news is that RCOs may know or discover that the control is already implemented (or partially implemented), especially where the RCO inherits enterprise controls from a parent organization with an established cybersecurity program. If implementation is partial, document how. We strongly recommend documenting the sources of information consulted in determining the current state. A well-written description of the current state will assist decision makers in determining next steps.
- (c) **Determine actions.** Once the current status for a control is determined, decision makers with risk acceptance responsibility (*see Must 6*) must determine whether the current status is acceptable or whether actions should be taken to enact change. For example, a decision maker may find that a particular control is relevant, but not currently implemented at all. It is the risk acceptor’s responsibility to determine what action, if any, should be taken. In documenting actions to be taken, we recommend that RCOs identify next (or first) steps in addition to the desirable end state. In this decision making phase, all 4 Pillars come into play, and RCOs must closely consider **Must 16 (Additional & Alternate Controls)** when making resource (re)allocation decisions. In some communities, comprehensive use of a general baseline control set may mitigate all but a small portion of cybersecurity risk. This may not be the case for all RCOs,¹⁸⁹ and adequate resourcing to mitigate unacceptable risks may pull significant resources toward additional and alternate controls.

RCOs should consider the following tools:

-  Trusted CI maintains a **CIS Controls Tracking Tool** template spreadsheet designed to support this baselining process for the CIS Controls.¹⁹⁰ Trusted CI designed this spreadsheet to capture an as-simple-as-possible layout for baselining against this control set.

¹⁸⁹ E.g., non-commodity assets and mission support often require greater emphasis on data and system integrity and availability than on confidentiality.

¹⁹⁰ <https://docs.google.com/spreadsheets/d/1XozhP8QY9mdm1nQyY5YOC26SexgDCqaakBYS4KVcEBg>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

🔑 Trusted CI maintains several templates which can be used to describe how controls are applied to an organization’s assets or classes of assets.¹⁹¹ In particular, the **Information Asset Inventory Template** and **Asset-Specific Access and Privilege Specification (ASAPS) Template** provide a starting point for capturing the way in which control baselines are applied to identified asset classes. The specific policy template such as the **Physical Security Policy Template**, **Password Policy Template**, **Asset Management Policy Template**, and **Social Media Policy Template** can be tailored to capture an organization’s baselines for specific types of controls found in many control sets.

We encourage RCOs to manage the control implementation activities using their usual project management methods. Good programmatic and project management involves prioritization of effort, and the implementation of many controls will amount to projects of some complexity and length (as opposed to discrete, quick turnaround assignments).

Common Challenges & Recommendations

This section describes some common **Must 15** challenges and offers recommendations on how to overcome them.

→ **Kickstart that program!**

An existing organization may need to kickstart its program quickly, implementing a baseline control set that allows the organization to mitigate risks to the most valuable and sensitive assets in a hurry. A newly funded organization may have the outlook and time to lay out a comprehensive vision of its cybersecurity program from a more top-down perspective: focusing on fleshing out governance and policy, considering how to integrate cybersecurity into broader processes, and making architectural decisions that will build security into processes from the start. That is ideal but not always realistic. Trusted CI recommends the CIS Controls as the starting point in both situations.

The CIS Controls Version 7.1 includes guidance in the form of “implementation groups” designed to simplify the tailoring of security controls according to an organization’s scale. Implementation groups for each control are constructed from activities involved in implementing a security control. These collections are organized into the three implementation groups for each control according to difficulty and organizational complexity. This provides an effective first-pass starting point for producing an initial control baseline tailored to an organization’s mission. An organization should choose an implementation group appropriate for its scale and use these activities to determine which activities should be prioritized in order to begin to effectively implement security controls.¹⁹² The Trusted CI CIS Controls Version 7.1 Tracking Tool is a spreadsheet which can be used to document the applicability of these controls and their implementation status in an environment.

→ **Breaking the bank.**

Organizations faced with implementing a baseline control set may be immediately deterred by the

¹⁹¹ <https://trustedci.org/framework>.

¹⁹² Note that the Implementation Group 3 of the CIS Controls is designed for organizations larger than typical RCOs. *Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0*

cost, either due to the sheer number of controls required, the complexity of implementing those controls, or the impact those controls will have on the mission. However, the Trusted CI Framework provides organizations with several tools to help mitigate the potential cost of full adoption of a control set. **Must 16 (Additional & Alternate Controls)** provides organizations the flexibility to select alternate controls that are more affordable or which have a less significant impact on their mission. **Must 6 (Risk Acceptance)** empowers organizations to accept the risk that comes from not implementing a control that is deemed too costly.

Moreover, the decision to accept the risk from not implementing a specific control or controls is not a permanent decision, and organizations can choose to defer implementation of less critical controls, accepting the risk arising from that delayed implementation. This allows organizations to amortize the cost of implementing their baseline control set over a longer period of time, thus making the baseline control set more affordable.

→ Under pressure.

Many RCOs are under pressure to adopt compliance-driven control sets, for instance NIST SP 800-53 and SP 800-171, due to contractual or legal requirements. RCOs not under an imminent threat of regulations (that mandate the NIST control sets) can take comfort in the knowledge that the CIS controls overlap significantly with NIST and other control sets. It provides a solid foundation for compliance if it becomes necessary in the future. Opting into NIST SP 800-53 or 800-171 as a general purpose baseline control set should be done only with serious consideration of the cost of implementation, the subset of controls (low, moderate, or high) that are most appropriate for the organization's mission in the case of 800-53, and the confidentiality-focus of the controls selected for 800-171.¹⁹³

¹⁹³ For a detailed discussion of the original purpose of NIST SP 800-171 and its limitations as a general purpose baseline control set, see <https://blog.trustedci.org/2017/06/nist-sp-800-171-and-its-potential.html>.



Must 16: Additional & Alternate Controls

Organizations must select and deploy additional and alternate controls as warranted.

Controls are specific administrative, technical, and physical safeguards and countermeasures. The specific controls included in baseline control sets may be insufficient in total to optimally balance risk mitigation with risk-taking necessary for mission success. Additional controls are those deployed to address unacceptable risks not covered by the baseline. Alternate controls are those deployed to mitigate unacceptable risks if implementing the alternate controls has a more positive impact on mission success than the baseline control.

Research organizations have a variety of traditional and specialized assets. This includes, but is not limited to, servers, desktops, supercomputers, research networks, instruments, sensors, data, and software that controls instruments. Baseline control sets are designed to address the most common risks in traditional IT environments and are not intended to cover such a wide range. Organizations also have a unique mission and culture, budget challenges, and existing workflows that constrain the implementation of baseline controls. Identifying these assets and constraints is key to designing additional and alternate controls.

Trusted CI strongly recommends that research cyberinfrastructure operators (RCOs) maintain awareness of community-driven and specialized approaches to additional and alternate controls.

Why is this a Must?¹⁹⁴

The primary purpose of additional and alternate controls is to allow organizations to meet the specific needs of the organization's mission. Baseline control sets are powerful tools, but they are also limited, and they are not designed to address the nuanced or unique requirements of every organization that adopts them. Rather, baseline control sets establish a default, which is supplemented by an organization via additional and alternate controls to suit the needs of its mission.

Although additional and alternate controls¹⁹⁵ both exist to allow an organization to tailor its cybersecurity strategy, they do so in different ways.

¹⁹⁴ The reader may note that this Must is substantially longer than others, which is by design. It is meant as a catalog that attempts to capture alternate and additional controls as comprehensively as possible.

¹⁹⁵ Note: Although additional and alternate controls are conceptually distinct, there are cases where the distinction can be murky. For instance, a baseline control set may under-defend against a specific threat the organization faces, while also over-defending against a threat the organization doesn't face. In this scenario, the organization could view its decision as replacing an unnecessary control with an alternate control that addresses a specific risk, or as selecting an additional control and accepting the risk from not implementing the unnecessary control. In either case, the important element is that the organization is tailoring the baseline control set to suit the needs of its mission.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

Additional controls allow an organization to **supplement** its baseline control set implementation when the level of protection it provides is insufficient to meet the needs of the organization's mission. This may be because the organization performs uniquely critical functions, because it faces unique threats, because it operates specialized systems that aren't sufficiently covered by the baseline control set, or because of the organization's higher or lower risk tolerance.

Unlike additional controls, alternate controls allow for organizations to **modify** what the baseline prescribes when the baseline controls are not compatible with the organization's mission. This may be because the baseline control set prescribes controls that restrict or inhibit the organization from performing its mission, because the controls would interfere with established organizational workflows, because of organizational resource constraints, or because the organization can achieve acceptable results more efficiently.

Additional and alternate controls are particularly important for RCOs, as RCOs tend to have specialized assets and unique missions which are not under the intended purview of most baseline control sets:

1. RCOs frequently prioritize information asset integrity and availability over confidentiality;
2. RCOs frequently build and operate customized technologies;
3. RCOs frequently prioritize collaboration and permissive access to information; and
4. RCOs serve a wide range of research missions and diverse collaborators such as international domain scientists.¹⁹⁶

As such, although baseline control sets are still valuable tools for RCOs, they are likely to require greater-than-average supplementation with additional and alternate controls to make the implemented control set a good fit for the organizational mission. The baseline control set RCOs select in **Must 15 (Baseline Control Set)** will have a significant impact on their implementation of **Must 16**, as selecting an in-depth baseline control set will require a number of alternate controls to be compatible with the organization's mission;¹⁹⁷ whereas selecting a minimal baseline control set will require significant additional controls to bring the organization's total risk down to acceptable levels.¹⁹⁸ As the Center for Internet Security (CIS)¹⁹⁹ points out, "You must still understand what is critical to your business, data, systems, networks, and infrastructures, and you must consider the adversarial actions that could impact your ability to be successful in the business or operation."²⁰⁰

The Roadmap

This section describes the steps needed for an RCO to supplement the baseline and devise or explore alternate controls in concert with thoughtfully using a baseline control set.²⁰¹ Step 1 is about identifying when and where alternate and/or additional controls are needed. Step 2 outlines

¹⁹⁶ Which may involve additional controls to respond to, *e.g.*, export control.

¹⁹⁷ *E.g.*, FISMA Moderate.

¹⁹⁸ *E.g.*, Australian Essential Eight.

¹⁹⁹ <https://www.cisecurity.org/>.

²⁰⁰ CIS Controls Version 7.1, pg 7.

²⁰¹ See also reading on tailoring and alternative/additional controls -

<https://www.sciencedirect.com/topics/computer-science/security-control-baseline>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

For Public Distribution

strategies for choosing alternate controls. Step 3 is about designing additional controls when the baseline comes up short.

→ Step 1. Determine need.

Additional and alternate controls are needed where the baseline controls are inadequate or negatively impact the mission in some way, for instance when they are deficient in protecting certain assets or when they create unacceptable impediments to productive research. Factors that drive additional and alternate controls can range from simple to complex or tricky, as illustrated below.

Specialized Assets

Technologies that are outside the mainstream of traditional IT present special challenges in cybersecurity. It is not uncommon for a project or organization to need to determine how to secure mission critical systems for which there is little or no explicit cybersecurity guidance available. In addition to commodity off-the-shelf hardware and software, many projects RCOs support unique, atypical, or custom-built equipment that serves a specific role in the research process. For instance, interferometers may be used for measuring gravitational waves and Digital Optical Modules (DOMs) for measuring neutrinos. Projects measuring environmental data across a large geographic area may use specialized field instrumentation to create wireless sensor networks.

Specialized equipment requires the same security attention as commodity hardware and software and often even greater attention and additional controls, as such equipment may constitute an organization's largest investment and most mission critical assets. Sometimes specialized equipment is controlled through a remote network connection that increases the risk of unauthorized access and/or abuse. The abuse or misuse of certain types of equipment can result in physical damage to an instrument and its surroundings, including a threat to human life.

Must 3 (Information Assets), Must 4 (Asset Classification), and Must 15 (Baseline Control Set) contribute to the identification of specialized assets, but technologies and equipment most likely to require additional and alternate controls include, but are not limited to, the following:

- Custom/specialized equipment/instruments
- Industrial Control Systems (ICS)/Supervisory Control and Data Acquisition (SCADA) systems that drive instruments such as a telescopes or a particle accelerators
- Field equipment such as balloons, rockets, and sensors
- Internet of Things (IoT) devices
- Medical devices
- Spacecraft
- Automobiles
- Drones and aircraft
- Custom software

Information

Information is an important organizational asset but is not typically viewed separately from the

information systems where it resides. The two get lumped together because system controls often dominate information security. However, information can drive the need for additional and alternate controls in and of itself. For instance, information subject to rules and regulations such as sensitive defense information, data on controversial topics such as climate science, or data with a high market value such as intellectual property can make a system attractive to attackers.

Workflows

A workflow can introduce its own risk apart from the aggregate risk of the assets comprising it and drive the need for additional and alternate controls. Consider, for instance, a researcher who is using a “full disk encrypted” laptop and a secure storage server but fails to encrypt sensitive data prior to transferring it to a third party, believing it remains encrypted. Mitigating this risk requires user education, not technical controls. Another, particularly challenging example is a workflow that completely bypasses organizational security, so-called “shadow IT.”²⁰² Anticipating and mitigating this risk requires documenting not just assets, but the reasons that drive the behavior and adding resources and controls to address them.

A different area where additional and alternate controls are required is when baseline controls can disrupt an existing workflow, for instance, removing administrator privileges across the organization, or blocking outside access to a system needed by a researcher while traveling.

Collaboration

RCOs are often subject to other constraints that drive the need for additional and alternate controls, for example, when a prescribed baseline control is inconsistent with collaboration, efficiency, or functionality. For instance, protectively isolating the work of each research project prevents or inhibits collaboration and synergy that may have been possible between the projects.

→ Step 2. Devise/Explore alternatives.

An alternate control²⁰³ refers to a control that produces the same or lesser risk or threat mitigation as another control (*e.g.*, a control in a baseline control set). In practice, constraints and risk tolerance dictate how an alternate control can or should be designed, especially for RCOs. To help them, we recommend the following strategies for alternate controls.

1. Implement an alternate control that mitigates risk more effectively than the original control, for example, an air gap for a critical ICS and SCADA system instead of firewalling it.
2. Choose a control that reduces the risk equally well. For instance, instead of CIS Controls Version 7.1 Control 5.2 (Maintain secure images or templates for all systems in the enterprise based on the organization’s approved configuration standards), the organization might use documented procedures to install and configure a system. Or, the organization may opt to

²⁰² According to Cisco (<https://www.cisco.com/c/en/us/products/security/what-is-shadow-it.html>), shadow IT is “... the use of IT-related hardware or software by a department or individual without the knowledge of the IT or security group within the organization.” It is often driven by security controls disrupting workflows or the organization failing to provide a secure solution.

²⁰³ We object to the sometimes-used term “compensating control” due to its negative denotation.

implement a Science DMZ²⁰⁴ when a firewall would negatively impact high volume data flows.

3. Use a control that reduces the risk partially and accept the residual risk. For example, use regular vulnerability scanning in lieu of a costly, full penetration test or full disk encryption instead of file level encryption to avoid managing encryption keys.
4. Leverage an external service provider who can address the risk equally well or better. For example, implement the functionality within a secure cloud environment.

In choosing options, consider factors such as cost, expertise, and level of effort needed to implement the control, compatibility with organizational culture and mission, risk tolerance, rules and regulations, and sustainability.

Finally, and most importantly, keep the following in mind while devising alternate controls:

1. **Think risk outcomes;** don't get tangled up in controls. Any approach that produces nearly the same outcome counts as a good alternative.
2. **Don't narrow your scope.** A certain type of control need not replace the same type. For example, a technical control can be replaced by an administrative control such as a policy with logging and reporting or a physical control such as restricting physical access.
3. **Don't reinvent the wheel.** Search for existing alternatives. Sometimes, alternate controls such as a Science DMZ architecture are readily available and well documented. Talk to peers that have similar assets and environment. Scour other control sets.

→ Step 3. Supplement.

Adding controls to a baseline to address specific risks can range from being relatively straightforward, for example by leveraging a different control set, to extremely challenging if no explicit security guidelines are available, *e.g.*, for specialized assets. As in devising alternate controls, constraints, and risk tolerance will again dictate the choice of additional controls, but unlike alternate controls, additional controls are not constrained by the baseline.

For RCOs needing guidance on supplementing baseline controls, we recommend the following:

1. **Break it down.** Many unique-looking information systems are aggregates of well-understood techniques and technologies. Often, breaking something complex into component technologies (*e.g.*, web application/portal, database servers, and job queues) will produce useful insight as to what controls may be needed (bottom up). Another, often complementary, approach is to carefully characterize the whole system, and work from there (top down).
2. **Don't reinvent the wheel.** Talk to peers and search for and leverage existing controls or guidance. For example, if CIS controls are used for your current baseline, look to other baseline control sets such as NIST 800-53 or ISO 27002.²⁰⁵ For specialized asset categories or areas, explore if a control set or guidance is available already, as it is for ICS and SCADA

²⁰⁴ Science DMZs can also be an additional control, depending on circumstances.

²⁰⁵ <https://www.iso.org/standard/54533.html>.

systems (*see* next section), medical devices,²⁰⁶ IoT devices,²⁰⁷ automobiles,²⁰⁸ spacecraft,²⁰⁹ supply chain security,²¹⁰ and cloud computing.²¹¹

3. **Focus on outcomes.** Similar to the approach for alternate controls, focus on outcomes instead of controls, and use administrative controls to replace technical controls where it makes sense.

Common Challenges & Recommendations

This section describes some common **Must 16** challenges and offers recommendations on how to overcome them.

→ Calculating cost.

Calculating the cost of an additional or alternate control can be a complex undertaking. Selecting a control to mitigate a particular risk is not an isolated event. A control might have side effects or bring along other vulnerabilities or restrictions, so total costs must be evaluated. For example, a decision to implement a backup system may in turn require ensuring the backups are encrypted, and that, in turn, requires the encryption key(s) be stored in a secure location but still accessible in foreseeable circumstances when the data needs to be restored.

→ Administrative or technical?

Deciding when and how to replace a technical control with an administrative control²¹² can be challenging for IT personnel with a strong technical focus. The rule of thumb is that an administrative control, especially with logging and reporting, can be as effective as a technical control, or a strongly enforced sanction policy.

→ ICS and SCADA system security.

ICS and SCADA systems present unique challenges due to a combination of mission criticality, the uniqueness of individual ICS and SCADA systems, especially in scientific contexts where systems may have a longer lifetime than the companies that produced them, and the difficulty of maintaining the security of ICS and SCADA components. CIS provides a companion guide for industrial control environments.²¹³ The Industrial Control System Cyber Emergency Response Team (ICS-CERT)²¹⁴ is a recognized resource in this area. NIST also provides guidance²¹⁵ on ICS security. Moreover, the

²⁰⁶ <https://www.fda.gov/media/86174/download>.

²⁰⁷ <https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259-draft2.pdf>.

²⁰⁸ https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/812333_cybersecurityformodernvehicles.pdf.

²⁰⁹ https://aerospace.org/sites/default/files/2019-11/Bailey_DefendingSpacecraft_11052019.pdf.

²¹⁰ <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-161.pdf>.

²¹¹ <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-144.pdf>.

²¹² A policy, procedure, or training.

²¹³ <https://www.cisecurity.org/white-papers/cis-controls-implementation-guide-for-industrial-control-systems/>.

²¹⁴ <https://ics-cert.us-cert.gov/>. *See, also*, NIST SP 800-82r2. Available at

<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf>.

²¹⁵ <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

For Public Distribution

security of these devices has much in common with Internet of Things (IoT) devices for which the Department of Homeland Security (DHS) has published a set of strategic security principles.²¹⁶

Special Topics

→ Identity management for distributed science communities.

Authentication and other identity management functions (including front-end interfaces) for a diverse population of end users are difficult to implement securely and expensive to maintain. Significant opportunities exist to outsource some of the issues around distributed identity management²¹⁷ as an alternate control. Organizations participating in a distributed identity management framework should consider joining the Sirtfi trust community²¹⁸ to facilitate incident response collaboration in a federated identity environment.

→ Long-term analysis and data retention.

Data collected by research projects is often of interest long after collection has ceased and funding has expired. NSF proposals generally require additional controls such as a data management plan that includes archiving the data with appropriate security protection.²¹⁹ In some cases, the ability to perform data analysis requires the preservation of a computing environment including stable versions of operating systems, analysis software, and associated software libraries. This environment may be implemented in as an internal cloud²²⁰ depending on security and performance requirements. Care must be taken to control access to unsupported systems and software and limit the impact of potential compromises of vulnerable systems. RCOs should consider providing facilities as part of their cyberinfrastructure to simplify the ability of research projects to satisfy long term processing and access requirements.

→ Securing scientific data and data flows.

Scientific data flows, while often not a confidentiality concern, require special attention to maintain integrity and availability. They require additional controls such as end-to-end network design for devices producing scientific data to ensure adequate integrity, performance, and reliability. Also, sufficient local storage may be needed as an additional control to allow scientific data collection

²¹⁶

https://www.dhs.gov/sites/default/files/publications/Strategic_Principles_for_Securing_the_Internet_of_Things-2016-1115-FINAL....pdf.

²¹⁷ <https://www.incommon.org/federation>.

²¹⁸ <https://refeds.org/sirtfi>.

²¹⁹ https://www.nsf.gov/pubs/policydocs/pappg19_1/pappg_2.jsp#IIC2j.

²²⁰ See BaBar Status and Data Preservation

<https://indico.cern.ch/event/588219/contributions/2371330/attachments/1427230/2190312/BaBar-DPHEP-20170314.pdf>, Design, Status, and Experience with Babar LTDA

<https://indico.cern.ch/event/209688/contributions/1501412/attachments/326047/454750/LTDA-DPHEP-Nov2012.pdf>, and the wiki located at

<https://confluence.slac.stanford.edu/display/BBRLTDA/The+BaBar+Long+Term+Data+Access+Archival+System+Project>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

For Public Distribution

during network outages. As previously mentioned, the Open Science Cybersecurity Risk Profile (OSCRP) is a potential resource for addressing these security concerns. The Trustworthy Data Working Group's (TDWG) "Guidance for Trustworthy Data Management in Science Projects"²²¹ is another available resource on this topic. The TDWG report showcases data owner and maintainer expressed needs, and identifies barriers to trustworthiness as well as tools and technologies for trustworthy data.

→ Non-RCO device access to RCO networks and data.

Visiting scientists and even employees may expect internal network access using their own desktops, tablets, or mobile devices; however, these devices might be unmaintained or be configured for a different environment. Network and service design may need additional controls to reduce the risk of unauthorized transfer of sensitive data or compromise of mission critical assets. Typically, this design segments non-RCO-managed devices to specific network segments and restricts access from those devices to specific services or service intermediaries like bastion hosts or virtual desktops. Administrative such as policies and procedures or technical controls may be employed to prevent unmanaged devices on networks with managed devices such as guest network/WiFi access and 802.11x authentication

→ Physical and environmental security.

RCOs may rely on physical and environmental security controls to maintain the availability of equipment and instruments and to ensure the integrity of scientific data. However, they also may face the need to give partial access to the public, manage data flows between geographically disparate sites, or operate in extreme environments. Additional or alternate controls can include door locks, gates, fencing, monitored security cameras, backup electricity generation, fire protection (with appropriate cutoffs), HVAC, and hardened or redundant configurations for hardware and software.²²²

→ Secure application software.

RCOs should identify all applications in the information asset inventory and prioritize additional or alternate controls needed to secure them by risk to the organization's mission. Moreover, organizations could consider training developers as an additional control to utilize secure coding techniques and tools as appropriate for the application implementation language(s).²²³ Developers of web applications can also leverage available guidance²²⁴ on how to secure them. Application developers should understand their responsibility to produce secure applications and work cooperatively with the cybersecurity team to uncover and remediate flaws as early as possible in the software life cycle.

²²¹ <https://doi.org/10.5281/zenodo.4056241>.

²²² <https://www.sans.org/reading-room/whitepapers/physical/physical-security-important-37120>.

²²³ <https://social.technet.microsoft.com/wiki/contents/articles/7100.the-security-development-lifecycle.aspx>.

²²⁴ The OWASP Top 10[#] (https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project) and the Securing Web Application Technologies (SWAT) checklist (<https://software-security.sans.org/resources/swat>).

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

→ Science gateways.

An alternate or additional control that may shield cyberinfrastructure from direct attacks or inexperienced users is a science gateway. It provides a set of tools, applications, and data collections that are integrated via a portal or a suite of applications to serve the needs of a specific scientific discipline, for example astrophysics, genomics, etc. Gateways provide access to a variety of capabilities including workflows, visualization, resource discovery, and job execution services, allowing scientists around the world to more effectively collaborate using shared cyberinfrastructure and data. More information may be found at the Cybersecurity for the Modern Science Gateways page.²²⁵

→ Science DMZ.

Another additional or alternate control in a research computing environment is a Science DMZ.²²⁶ It is an approach to provide the low-latency, high-bandwidth connections demanded by scientific computing applications without giving up on network security. Because a Science DMZ exists at the edge of its local network (or as close as possible thereto), typically outside institutional firewall and monitoring infrastructure, seeing to security needs typically falls on the RCO or central IT running the DMZ rather than larger institutional resources.

Much of a Science DMZ's speed is owed to its avoidance of stateful firewalls and traffic shaping mechanisms normally in place on large enterprise networks. Network-level security is often handled largely or even solely through router ACLs. This particular approach is most effective when those ACLs are automatically updated in response to information gathered by an intrusion detection system (IDS) such as the Zeek network security monitor.²²⁷ Zeek keeps detailed records of application layer state. Its analytics can be mapped to behavioral patterns of end users, so automated security scripts can respond in real time, *e.g.*, updating router ACLs or notifying security on-call personnel.

In addition to helping to keep traffic moving, network performance monitoring tools²²⁸ can enhance security on a Science DMZ. Such tools will give a clear understanding of how the flow of traffic is moving on the network, information that may aid in coping with, for example, Distributed Denial of Service attacks.

Further reading: Science DMZ Security - Firewalls vs. Router ACLs.²²⁹

→ Situational Awareness.

An additional control to enhance cybersecurity is collecting session data.²³⁰ It is useful because it can be done passively, it provides situational awareness for the network without adversely affecting

²²⁵ <https://sciencegateways.org/-/webinar-cybersecurity-for-the-modern-science-gateway>.

²²⁶ <https://fasterdata.es.net/science-dmz/>.

²²⁷ <https://zeek.org/>.

²²⁸ Such as perfSONAR, <http://www.perfsonar.net/>.

²²⁹ <https://fasterdata.es.net/science-dmz/science-dmz-security/>.

²³⁰ Such as netflow or IPFIX.

network performance. This data, especially when correlated with data from an IDS like Zeek, can help identify outlier behaviors that may be early indicators of malicious activity, *e.g.*, an attacker's lateral movement within the network, command and control traffic, or data exfiltration. Since session data does not contain the contents of the traffic, it does not take up much space, allowing an extensive history to be saved and analysed. Session data suites²³¹ may be used for this type of data collection and analysis.

²³¹ For instance SILK (<https://tools.netsa.cert.org/silk/>) and Argus (<https://openargus.org/>).

Glossary of Key Terms

Access - logical or physical ability to view, create, modify, or destroy information, or modify or destroy information assets.

Additional Controls - controls deployed to address unacceptable risks not covered by the baseline control set.

Alternate Controls - controls deployed to mitigate unacceptable risks if implementing the alternate controls has a more positive impact on mission success than the baseline control.

Authority - legal, administrative, logical, or physical control of information assets.

Baseline Control Set - a predetermined set of controls used as a default when selecting security controls for information assets. The baseline control set does not determine what security controls an organization must implement; rather, it provides a foundation from which an organization tailors control selection based on the needs of its mission.

Controls - specific administrative, technical, and physical safeguards and countermeasures applied to reduce cybersecurity risk.²³²

Cybersecurity - “Prevention of damage to, protection of, and restoration of computers, electronic communication systems, electronic communication services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and non-repudiation.”²³³ This definition is scoped to include information assets beyond traditional information technology (IT), and includes operational technology. Generally, Trusted CI uses the terms “cybersecurity” and “information security” interchangeably.

Cybersecurity Budgets - financial plans that commit specific resources for the organization’s cybersecurity efforts over a designated period of time.

Cybersecurity Lead - the organizational role with primary responsibility to advise and provide services to the organization on cybersecurity matters. This position is often titled “Chief Information Security Officer.”

Cybersecurity Program - a group of related cybersecurity-focused projects and ongoing activities managed in a coordinated way to obtain benefits not available from managing them individually.²³⁴

Entities - individuals or organizations.

²³² Adapted in part from The Information Security Practice Principles Foundational Whitepaper, <https://cacr.iu.edu/principles/ISPP-Foundational-Whitepaper-2017.pdf>.

²³³ <https://fas.org/irp/offdocs/nspd/nspd-54.pdf>.

²³⁴ Adapted in part from Schwalbe, Information Technology Project Management, 9th Edition. *See also*, <https://www.pmi.org/learning/library/understanding-difference-programs-versus-projects-6896>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

External Resources - services, tools, and collaborators outside of the organization that can be leveraged to support the cybersecurity program.

Information Assets - valuable, sensitive, and/or mission critical information and information systems.

Information Asset Documentation - the collection of artifacts describing the cybersecurity relevant details of information assets presented in a form that is useful to cybersecurity professionals and decision makers.

Mission - 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.

Obligations - internally or externally imposed processes or practices that impact the operation of the organization's cybersecurity program.

Organizational Leadership - Senior executives and other decision makers responsible for an organization. These are the people ultimately responsible for the organization who make final decisions regarding the highest priorities. Common leadership roles/titles include Director, Board, Chairman, Chief, Executive, Commander, President, Vice President, Partner, Principal, Owner, Founder, and Secretary.

Operational Technology - "hardware and software that detects or causes a change, through the direct monitoring and/or control of industrial equipment, assets, processes and events."²³⁵

Personnel Resources - 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.

Policy - 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.

Programmatic Evaluations - how the organization determines whether the cybersecurity program is achieving its purpose.

Programmatic Refinements - changes designed to improve the program's efficiency or effectiveness. Evaluation and refinement of a cybersecurity program can take many forms depending on the formality and scope of the assessment and the type of evaluation (e.g., planned, comprehensive program evaluations; internal self-evaluations following an incident).

Research Cyberinfrastructure Operator (RCO) - Organizations that operate on-premises, cloud-based, or hybrid computational and data/information management systems, scientific instruments, visualization environments, networks, and/or other technologies that enable knowledge breakthroughs and discoveries. These include, but are not limited to, major research facilities (e.g.,

²³⁵ See <https://www.gartner.com/en/information-technology/glossary/operational-technology-ot>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

NSF Major Facilities, fka Large Facilities), research computing centers within research institutions, and major computational resources that support research computing.

Risk - uncertain events or conditions—such as a successful cyber attack—that, if they occur, have a positive or negative effect on the organization’s mission.

Risk Acceptance - a decision to acknowledge a risk and not take further action unless the risk occurs.

Stakeholders - people or entities with interest in or affected by an organization’s cybersecurity.

Appendix A - Does my Organization Need its own Cybersecurity Program?

Does my unit, virtual organization, or collaboration need a cybersecurity program of its own?

For stand-alone organizations, businesses, and legal entities (*e.g.*, a corporation, an LLC, a university, governments), there is no question: These organizations need to treat cybersecurity programmatically just like they treat safety, human resources, finance, and other major business units and functions programmatically. For even very small organizations (*e.g.*, a family-run hardware store), we recommend going through the Musts and asking, “Have we addressed this? If not, who has? Should we?”

However, the answer is not so simple when it comes to (a) organizations that are **units** of larger organizations (*e.g.*, a research institution or supercomputing center part of a large university) and (b) **virtual organizations** (VOs)²³⁶ and **collaborations** between or among multiple organizations.

Units. In the case of a unit of a larger organization, the question is not whether to handle cybersecurity independently of the parent organization, but whether a cybersecurity program at the unit level is warranted to protect both the unit’s (and parent’s) missions and interests. Having a cybersecurity program need not (and probably should not) be exclusionary of the parent organization’s cybersecurity services or policies, and it may be able to rely heavily on the parent organization.

Collaborations / Virtual Organizations. Similar to units, collaborations need not (and almost certainly should not) try to formalize a cybersecurity program to the exclusion of the various partners’ respective cybersecurity services and policies. Some collaborations may be able to rely largely if not entirely on the partners doing their parts, and agreements that govern the relationship. However, as seen frequently in the research community, collaborations, and (indeed) virtual organizations often have distinctive missions and needs. A cybersecurity program is worth considering.

Both units of larger organizations and collaborations/virtual organizations (collectively “organizations”) should consider the following factors in determining whether to formalize a cybersecurity program. “Yes” answers to any of these questions should add a weight on the scale in favor of formalizing a program.

- 1) Is the organization’s mission extremely important or risky? [**Must 1 (Mission Focus)** or **Must 9 (Policy)**]
- 2) Does the organization have a history of or anticipate that it may face cyber threats (adversaries and attack types) that are different from the parent? [**Must 9 (Policy)**]
- 3) Is the organization very large and complex on its own?
- 4) Does the organization have stakeholders and cybersecurity obligations that are distinctive

²³⁶ See for examples of virtual organizations: <https://opensciencegrid.org/about/organization/>.

Trusted CI Framework Implementation Guide for Research Cyberinfrastructure Operators v1.0

compared to the major stakeholders and obligations of the parent? [**Must 2 (Stakeholders & Obligations)**]

- 5) Does the organization have a distinct set of users or suppliers significantly different than those of the parent organization? [**Must 2 (Stakeholders & Obligations)** and **Must 8 (Comprehensive Application)**]
- 6) Does the organization operate information assets that are different in meaningful ways from those of the parent? [**Must 3 (Information Assets)** and **Must 16 (Additional & Alternate Controls)**]
- 7) Does the unit have leadership roles with significant autonomy or discretion in terms of risk taking, budget, hiring, business development, and/or procurement? [**Must 5 (Leadership)**]
- 8) Does the parent's baseline control set and implementation clash with the unit's mission? [**Must 1 (Mission Focus)** **Must 15 (Baseline Control Set)** and **Must 16 (Additional & Alternate Controls)**]
- 9) Is the unit's mission highly distinctive in some other way that warrants special attention and may be outside the standard operations for the majority of the rest of the business? [**Must 1 (Mission Focus)**]
- 10) For collaborations and virtual organizations specifically. Does the collaboration have a user base that is shared across the partners (*e.g.*, cross partner scientific workflows with identity management and incident response implications)? [**Must 8 (Comprehensive Application)**]

Appendix B - Trigger Events

	Must 1	Must 2	Must 3	Must 4	Must 5	Must 6	Must 7	Must 8	Must 9	Must 10	Resource Musts	Must 15	Must 16
Changes in the organizational mission	X				X						X		
Reorganization of unit responsibilities	X				X	X							
Changes in stakeholder obligations including legal or government regulations		X						X			X		
Initiating relationships with new stakeholders		X						X			X		
Oversight or funding agency review of program milestones, metrics, and incidents		X								X			
Research projects using the CI that come and go		X						X			X		
Projects involving new IT assets, especially unique assets		X	X	X							X	X	X
Software upgrades (OS and major application)			X									X	
Hardware additions, subtractions, and upgrades (IT and OT)			X	X							X	X	X
Additional internet-accessible systems			X			X					X	X	
Changed or additional categories that alter the control requirements				X							X	X	X
Changes to risk acceptance role delegation						X					(11)		
Addition of new units managing risk acceptance						X					(11)		
Changes to cybersecurity lead personnel or reporting relationship							X						
Change to contracts, MOUs, or other agreements		X						X			X	X	
Initiating relationships with new entities		X	X	X				X			X	X	
Policy, procedure, or control changes								X	X		X	X	X
Changes in risk acceptance and determination of adequate cybersecurity resources					X	X	X				(11)	X	
Changes in budget or funding structure					X						X	X	
Changes in resources available to the program					X		X				X	X	
Changes in external resources, services, and technology available					X		X		X		X	X	X
Changes in the organization's attack surface (new services, ports)			X	X					X		X	X	X
Updates to baseline controls due to cybersecurity incidents and breaches				X							X	X	
Updates to baseline controls due to industry or other compliance requirements				X							X	X	
A significant change in threats or attack tactics					X		X		X		X	X	X
New technologies implemented for defense improvement.							X				X	X	X
Likely attacks that increase in both ease-of-use or sophistication.							X				X	X	X
Incidents resulting in unauthorized access to or exfiltration of data			X	X					X		X	X	X

Appendix C: Baseline Controls Sets

This appendix provides an overview of some of the best, commonly used, or commonly referenced baseline controls sets in the RCO community.

The Center for Information Security (CIS) Controls²³⁷ are widely regarded as an authoritative, reasonable, and prioritized set of controls.²³⁸ CIS also publishes the CIS benchmarks which provide recommended secure configuration guidelines for over 140 commonly used information technology products as well as additional resources including auditing tools and baseline images.²³⁹

The National Institute of Standards and Technology (NIST) Framework for Improving Critical Infrastructure Cybersecurity (“NIST Cybersecurity Framework” or CSF)²⁴⁰ was developed by NIST in collaboration with a range of private sector stakeholders for the protection of US critical infrastructure. NIST CSF is a voluntary framework. The “Framework Core” incorporates controls from a number of well known control sets and organizes them under five security “Functions”: Identify, Protect, Detect, Respond, and Recover.

NIST’s Risk Management Framework (RMF)²⁴¹ applies to federal government entities and some government contractors. The authorizing statute is the Federal Information System Modernization Act (FISMA).²⁴² Its baseline control sets (Low Moderate, and High) are encapsulated by NIST Special Publication 800-53, “Security and Privacy Controls for Information Systems and Organizations.”²⁴³

NIST Special Publication 800-171, “Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations” is a control set developed by NIST for use in federal contracts that involve the creation or handling of Controlled Unclassified Information (CUI).²⁴⁴ SP 800-171 is focused on protecting the confidentiality of CUI when it is handled by actors outside of the federal government. As such, the control set in SP 800-171 is derived from the FISMA Moderate baseline for confidentiality in NIST RMF, outlined in SP 800-53 (discussed above).

The HIPAA Security Rule²⁴⁵ is designed specifically to prevent unauthorized exposure of electronic protected health information (ePHI). It comprises a set of physical, administrative, and technical

²³⁷ <https://www.cisecurity.org/controls/>; for more detail including history, *see* <https://www.sans.org/critical-security-controls>.

²³⁸ The 2016 California Data Breach Report, <https://oag.ca.gov/sites/all/files/agweb/pdfs/dbr/2016-data-breach-report.pdf>, states that failure to implement the controls described by CIS that apply to an organization’s environment constitutes a lack of reasonable security. Further, it also states that multifactor authentication should be available for online accounts with access to sensitive information and that such information should be encrypted on portable devices.

²³⁹ <https://www.cisecurity.org/cis-benchmarks/>.

²⁴⁰ <https://www.nist.gov/cybersecurity-framework>.

²⁴¹ [https://csrc.nist.gov/projects/risk-management/risk-management-framework-\(RMF\)-Overview](https://csrc.nist.gov/projects/risk-management/risk-management-framework-(RMF)-Overview).

²⁴² <https://www.congress.gov/bill/113th-congress/senate-bill/2521>.

²⁴³ <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>.

²⁴⁴ <https://csrc.nist.gov/publications/detail/sp/800-171/rev-1/final>.

²⁴⁵ <https://www.hhs.gov/hipaa/for-professionals/security/index.html>.

safeguards that are statutory for healthcare-related organizations designated as “covered entities.”

ISO/IEC 27001²⁴⁶ is an information security standard published by the International Organization for Standardization²⁴⁷ and the International Electrotechnical Commission.²⁴⁸ It defines an Information Security Management System (ISMS) that is intended to bring information security under management control and gives specific requirements, including a control set.

The Australian Government identified through exhaustive research its mandatory “Top 4” Strategies to Mitigate Targeted Cyber Intrusions²⁴⁹ and later expanded its advice to the Essential Eight.²⁵⁰

The New Zealand Government has developed the Protective Security Requirements (PSR)²⁵¹ to outline the government’s expectations for security governance and for personnel, information, and physical security. The PSR mandatory requirements include a number of information security-related requirements in the Governance section²⁵² in addition to the requirements in the Information Security section.²⁵³

The Canadian government has produced guidelines for minimal cybersecurity controls for small and medium-sized organizations.²⁵⁴

Organizations using cloud services should investigate the CIS-provided hardened virtual images^{255,256} available for some of the major cloud service providers. Also, a review of the document “Security Best Practices for Academic Cloud Service Providers”²⁵⁷ is strongly recommended. Another good resource for cloud security is the Security Guidance²⁵⁸ from the Cloud Security Alliance (CSA).²⁵⁹

²⁴⁶ <https://www.iso.org/isoiec-27001-information-security.html>.

²⁴⁷ <https://www.iso.org/home.html>.

²⁴⁸ <https://www.iec.ch/>.

²⁴⁹ <https://www.asd.gov.au/infosec/top-mitigations/top-4-strategies-explained.htm>.

²⁵⁰ <https://www.asd.gov.au/publications/protect/essential-eight-explained.htm>.

²⁵¹ <https://www.protectivesecurity.govt.nz/>.

²⁵² <https://www.protectivesecurity.govt.nz/governance/>.

²⁵³ <https://www.protectivesecurity.govt.nz/information-security/>.

²⁵⁴ <https://cyber.gc.ca/sites/default/files/publications/Baseline%20Cyber%20Security%20Controls%20for%20Small%20a.nd%20Medium%20Organizations.pdf>.

²⁵⁵ <https://www.cisecurity.org/services/hardened-virtual-images/>.

²⁵⁶ <https://www.cisecurity.org/services/hardened-virtual-images/cis-hardened-images-faq/>.

²⁵⁷ <https://scholarworks.iu.edu/dspace/handle/2022/22123>.

²⁵⁸ <https://cloudsecurityalliance.org/research/guidance/>.

²⁵⁹ <https://cloudsecurityalliance.org/>.