



Whitepaper

The Five ICS Cybersecurity Critical Controls

Written by Robert M. Lee and Tim Conway

October 2022

Introduction

Five cybersecurity controls can be utilized together to create an efficient and effective industrial control system (ICS) or operational technology (OT) security program.¹ This paper identifies those controls and the rationale behind them so that organizations can adapt the controls to fit their environment and risks. The controls are intended to be outcome-focused instead of prescriptive in nature. They are also intelligence-driven in that they have been chosen based on the analysis of recent compromises and attacks in industrial companies around the world.

The five controls addressed in this paper are represented in Figure 1.



Figure 1. Five Critical Controls for ICS Cybersecurity

Organizations should note, especially if they are part of critical infrastructure, that they have an obligation to ensure a safe operating environment for personnel and a duty to protect from harm the communities they operate in by ensuring appropriate investments in ICS cybersecurity.

Organizations have no obligation, however, to exceed the minimums of mandatory or expected good practices to further protect their business interests. Said simply, organizations must participate in their own defense to protect community and national security against known threats, but they must make decisions according to risk tolerance and return on investment.

Organizations must participate in their own defense to protect community and national security against known threats, but they must make decisions according to risk tolerance and return on investment.

¹ For the purposes of this paper, ICS and OT are used interchangeably. The authors note that OT is the broader classification of systems whereas ICS is a more specific type of OT in industrial organizations. As an example, the building automation systems in a datacenter are OT systems but are not in an industrial environment whereas the automation system in a chemical plant would be ICS. Across the community, there is no one lexicon used but generally each means “not IT.”

The expectation that organizations will defend the operational assets to appropriate levels, make balanced risk decisions in the defense of business supporting digital assets, and have a firm grasp of all areas where these systems overlap or have interdependencies is unrealistic for leaders in complex organizations. As a result, business leaders across critical infrastructure sectors find themselves seeking guidance on which security controls and technology investments should be pursued programmatically, when in fact, the solutions they seek are more dynamic in nature. Threats and adversary actions change over time, the available technologies change with new innovations. Most importantly, the security maturity and capabilities of an organization change, informing programs that should be pursued at any given time. The five critical ICS security controls presented in this paper constitute what the authors assess to be the minimums for community and national security based on the real-world attacks they are derived from and designed to mitigate. Organizations can, of course, go beyond these five controls to further reduce risk according to the organizational goals and risk assessments.

Existing Frameworks and Guidance

Many excellent frameworks and guidance documents are available across the infrastructure community. NERC CIP for the North American Bulk Electric System,² ISA/IEC 62443,³ the US Department of Homeland Security's Cybersecurity and Infrastructure Security Agency's common performance goals, Cybersecurity Capability Maturity Model,⁴ Cybersecurity Maturity Model Certification,⁵ the National Institute of Standards and Technology's frameworks,⁶ and others all provide a wide view into the cybersecurity of industrial systems.

Framework Genesis and Purpose

Many of these frameworks were created when the insight into ICS-specific cyber threats was incredibly limited. In spite of this limitation, the drafting teams did an outstanding job of providing guidance and keeping the community updated. As a result of their origin, however, many of the security controls are often IT controls that can be applied to OT environments. This indirect applicability resulted in many controls that may be possible to apply now but do not provide a view into or reduce risk against existing cyber threats.

² Standards, NERC, www.nerc.com/pa/Stand/Pages/default.aspx

³ "New ISA/IEC 62443 standard specifies security capabilities for control system components," ISA, www.isa.org/intech-home/2018/september-october/departments/new-standard-specifies-security-capabilities-for-c

⁴ C2M2 Version 2.1, C2M2, <https://c2m2.doe.gov/>

⁵ "Cybersecurity Maturity Model Certification," Office of the Under Secretary of Defense, www.acq.osd.mil/cmmc/index.html

⁶ Cybersecurity Framework, NIST, www.nist.gov/cyberframework

Additionally, there is a prevention bias in the larger information security community. Among the most well-known and utilized frameworks, on average, between 60-95% of all the guidance is preventive in nature. The controls are designed to prevent compromises using approaches such as endpoint protection, hardening of devices, passwords, access management, and patching. As a result, many organizations invest as few as 5% of their resources to detecting, responding, operating through an attack, and recovering from compromises. Without robust detection and response insights, the preventive controls are limited and atrophy over time. Security teams and executives often believe their security is better than it is, due to this atrophy.

As regulations emerged for some sectors, many organizations chased projects to implement security controls. Other organizations shaped existing security programs to support compliance requirements.

The five ICS Critical Controls highlight the strength of an interdependent, balanced, preventative, detective, and responsive approach.

How to Utilize Frameworks

The five ICS Critical Controls highlight the strength of an interdependent, balanced, preventative, detective, and responsive approach. Using existing frameworks, they provide a common lexicon to communicate across the community and to non-security professionals. The controls also provide an opportunity for benchmarking against peer organizations, which is a useful tool at the board of directors and executive leadership level to help manage investments. The authors of this paper acknowledge the significant benefit of these frameworks and standards, and how they shape the cybersecurity and reliability of key resources over time. Those benefits include:

- **Prioritized rapid response recommendations.** During times of accelerated geopolitical events, leaders want to know what to focus on across the sea of regulations, frameworks, guidelines, and requirements. These five ICS Critical Controls are an answer to the frustrated organization request, “Just tell us what to do. We can’t do everything immediately, but where should we start?”
- **Getting more out of minimum requirements.** When no regulation exists, organizations invest in cybersecurity capabilities based on various risk management strategies. This strategy results in wild variations of approaches across a sector, with some organizations investing heavily in some technologies or capabilities, and others doing the absolute minimum. This disparity creates a cyber-target bell curve for adversaries to develop attack strategies with targets of opportunity on one end of the curve and targets of selection on the other.

Regulatory requirements are established as the minimum set of requirements that need to be achieved, but they suffer from regulatory process development and implementation lag, which means that it can take years before organizations implement and benefit from them. In some cases, requirement-based security programs create predictable common defense approaches across a sector that can be anticipated and targeted by adversaries.

For sectors with no regulation, the five ICS Critical Controls need to be a programmatic focus across prioritized operational assets. For organizations subject to larger regulation, these five ICS Critical Controls represent areas where efforts should be pursued to go beyond the baseline minimum requirements.

The authors recommend that organizations adopt the five Critical ICS Cybersecurity Controls and then map their organization's efforts into a common framework for communication around the organization and to peer organizations. That common lexicon will facilitate clear communication without introducing additional jargon. In fact, the first of the five controls asks organizations to identify the scenarios they want to be prepared to defend against; this is in keeping with guidance from various standards such as IEC-62443, which guides participants to conduct a risk assessment first instead of an unfocused test of every control.⁷

Key Considerations on the Differences Between OT and IT

Some key differences between OT and IT are relevant to discuss for the context of this paper. It is often cited that the differences between IT and OT include purpose-built systems, legacy systems, unique communications and network protocols, and the system's ability to adopt certain security controls. While often accurate, these considerations are a technology-focused view. In critical infrastructure, the biggest difference between IT and OT is the mission or business purpose of the systems. As a generalization, one (IT) is focused on *how* you manage the business while the other (OT) is focused on *why* you are a business. The mission, or purpose, of those systems dictates what is required of them and what the risks and threats are to those systems.

In critical infrastructure, the biggest difference between IT and OT is the mission or business purpose of the systems.

Common Technology, Unique Risks

A Windows operating system computer hosting a database for a financial institution has a distinctly different purpose and impact of failure when compared to a Windows operating system hosting the Human Machine Interface (HMI) for a nuclear power plant. An adversary may be able to exploit a targeted Windows system in a similar way but their behavior within that system will differ depending on whether they are focused on intellectual property theft of the financial institution's database versus causing an unsafe operating condition and physical impact in the nuclear power plant. This concept was covered in the SANS paper, "The ICS Cyber Kill Chain."⁸

⁷ Teams often argue not about the usefulness of a given security control but whether it effectively and efficiently addresses a specific risk. It is paramount for organizations to align first on what the risk scenarios are. This approach will help eliminate dogmatic approaches to standards where a control, such as software patching, is seen as a right or wrong action instead of a measure that may or may not make sense in the context of the operations.

⁸ The Industrial Control System Cyber Kill Chain," by Michael Assante and Robert M. Lee. SANS Institute, October 5, 2015. www.sans.org/white-papers/36297/ (Registration required to download.)

As an overly simplified abstraction, IT cybersecurity tends to be focused on system and data security. OT cybersecurity tends to be focused on system of systems and physics. Gaining access to the system and understanding the system or its data is critical to many IT compromises. The adversary does not often seek to cause physical manifestations in cyberattacks—the goals are more likely data theft or disabling of the systems. The types of attacks that most worry the OT cybersecurity community, though, are those that seek to disrupt operations, cause physical damage, or even cause safety-related incidents that reach the level of equipment damage or loss of life. It is improbable for an adversary to achieve such effects by targeting a single system. An adversary targeting OT for these types of outcomes needs to be able to take advantage of the system of systems and understand the physics of that environment. As an example, a compromised Engineering Workstation (EWS) in a well-designed environment is not usually enough to cause any significant issues in a production environment.⁹ To reasonably achieve a more devastating outcome, the adversary could compromise the EWS, learn how to manipulate the logic on a controller through the EWS, and with an understanding of the logic, impact the production process. In other words, the adversary can target System 1 to manipulate System 2 to cause a physical impact in System 3.

Understanding or overfocusing on any one system in OT misses the reality of typical adversary objectives. Because the OT environment, by design, must allow for this type of system-to-system interaction and functionality, most cyberattacks do not require the exploitation of vulnerabilities to achieve the adversary's goals. Instead, an adversary who learns the system of systems and physics can utilize required and native functionality to achieve their objectives. This is a radically different way of viewing cybersecurity than is present in most IT environments.

If the mission is different, the systems are different, the workforce skillsets are different, the adversary objectives and actions are different, the impact is different, and the organizational or business unit goals are different—then the security controls cannot be the same. IT and OT are different and at a minimum, the cybersecurity controls must be tailored and prioritized differently based on that reality.

Control No. 1: ICS-specific Incident Response Plan

Organizations must have an ICS-specific incident response plan to account for the complexities and operational necessities of responding in operational environments. A common mistake for organizations is thinking about incident response as a final element in its security program. This approach results in all the security controls and choices implemented earlier being misaligned with the necessities of incident response. As an example, organizations may find that the threat detection strategy, architecture choices, and data collection implemented does not support the requirements of the incident response. Additionally, with growing regulatory requirements on incident reporting worldwide, organizations must identify what questions and requirements need to be addressed for a successful incident response long before the incident occurs.

⁹ Adversaries learning environments and making changes can have unintended effects, including loss of control and safety. For effects such as physical damage to occur, however, the adversary not only needs confidence about what they want to achieve and how to achieve it, but also additional time to understand the unique operating environment and to develop the misuse and manipulation attack. The increased complexity and uniqueness of the target environment will also require additional resources with varied skillsets.

A key component of incident response in industrial organizations is the capability to gain root cause analysis. Root cause analysis supports the ability to return to safe operations, yet the growing complexity of industrial automation has made gaining root cause analysis more complicated. IT incident response plans often focus on identifying the adversary, containment, and eradication. OT incident response plans prioritize actions based on the potential for operational impact and how to position the system to operate through the attack in a manner that reduces the effect of the attack and impact on the process under control. Incident response and the cybersecurity investments to support it can not only reduce cyber risk but also enhance operational resilience because they facilitate root cause analysis of failure events, regardless of whether they are adversary-influenced.




Figure 2 presents an overview of the incident response planning process. As a first step, organizations should determine what scenarios they want to be able to reduce risk against and be ready to respond to in an adverse event.¹⁰ These scenarios create alignment from the board room to the plant floor, setting the requirements for everyone. Scenario planning is a natural occurrence for board members and executive leadership, is similar to scenarios and planning in the safety engineering community, and aligns to kill chain analysis, as is common for cybersecurity personnel. This commonly understood approach is not specific to any one skillset or profession, making understanding and buy-in easier. Cyberattacks are not singular events but part of an overall scenario, chain of events, or operation, so considering security controls in this way benefits defenders. Instead of trying to develop requirements against “virtual private network (VPN) compromises,” the organization can instead think through a full ransomware scenario.

Figure 2. ICS-specific Incident Response Planning Process

¹⁰ For information on scenario planning, please see www.weforum.org/agenda/2021/05/cybersecurity-safety-engineering/

Scenario Planning, Step 1

Scenario planning should begin with real-world incidents and be intelligence-driven. Organizations should determine what incidents have occurred in their industry and begin there. Some scenarios will be appropriate for multiple industries, whereas some may only be appropriate for specific types of environments in an industry. As an example, all organizations should have a ransomware scenario for ransomware in OT networks. An oil and gas company should have a safety-system-compromise scenario based on the TRISIS cyberattack but may opt only to apply those scenarios where there are safety systems in their organization.¹¹ Almost all power companies should have a Ukraine 2015 and Ukraine 2016 electric system cyberattack scenario. Yet, those companies may only think through those scenarios for their transmission and distribution operations rather than power-generation operations. However, organizations should not make the mistake of considering technology-specific factors such as the vendor of the safety system or protocol. As an example, the Ukraine 2016 cyberattack utilized the CRASHOVERRIDE/INDUSTROYER malware, which used IEC104 network protocols. A power company should not look at their operations, determine they use DNP3, and consider the scenario invalid. The point was the manipulation of circuit breakers and electric operations on the transmission system. The difference in protocol was based only on what the target environment contained, not the applicability of the operation.

Scenario Planning, Step 2

After considering the scenarios that have affected their industry (two or three intelligence-driven scenarios is a good starting place), the organization should consider a consequence-based scenario. Regardless of whether the attack has happened in the real world, identify a high-consequence impact that operations, engineering staff, or leadership are concerned about and map out whether it is achievable through a cyberattack. Chart out what the adversary would have to do to achieve this attack and use that scenario, as well.

Intelligence-driven scenarios should be prioritized because the likelihood of them being repeatable is high. Additionally, they are real-world scenarios that have already happened, because there are a small enough number of them, they can serve as especially useful focus areas. Consequence-based scenarios are “the art of the possible” and can become overwhelming. To prevent getting lost in the details and the fear, consider the scenarios within the context of your organization. Utilize the expertise and knowledge of the internal team, which has information that the adversary may not.

Consequence-based scenarios are “the art of the possible” and can become overwhelming. To prevent getting lost in the details and the fear, consider the scenarios within the context of your organization.

¹¹ “Triton/Trisis Attack Was More Widespread Than Publicly Known,” Dark Reading, January 16, 2019.
www.darkreading.com/attacks-breaches/triton-trisis-attack-was-more-widespread-than-publicly-known

Target two to three intelligence scenarios and one consequence scenario to guide the security program. For organizations in industries with no public cyberattacks to consider, utilize adjacent scenarios. As an example, mining operations have had limited visibility into cyberattacks but they also utilize safety systems; therefore, a TRISIS scenario adapted to what it would look like in their environments would be a great scenario to consider and model. Modelling adjacent industry scenarios for the organization is a useful way to be proactive beyond the “known unknowns” without utilizing resources against the “unknown unknowns” that may never manifest or manifest in ways that the security investments help against.

Tabletop It

Once the scenarios are chosen and agreed upon in the organization, the ICS-specific incident response plan should include a tabletop exercise (TTX). The TTX should overlay the scenarios against the organization’s environments and sites. Each part of the organization (such as operations, regulatory, security, business, and legal) should determine what requirements it would have for each incident. Determine what the organization will do in response to a given incident and what information it will need to know to inform actions. Identify these requirements up front and utilize the TTX to determine what will be needed from the environment and in what timeframe (for example, the accessibility of certain key datasets and their retention length articulated in a Collection Management Framework).¹² The findings from the TTX will help guide the organization about what sites are the most important to focus on, what the crown jewels are at those sites, and how to focus the remaining four critical controls.

Numerous sectors have strong events analysis programs and lessons-learned information-sharing forums (NERC E-ISAC, FAA, NRC, chemical safety, and pipeline safety, to name a few) that encourage and enable organizations to work through real-world scenarios and consider unique impacts. Many countries and sectors are also working on large-scale, joint team-training-focused exercises. These large-scale exercises go beyond individual organization level tabletop exercises as described above and highlight additional interdependencies across other sectors, supply chains, and governments.

An ICS-specific incident response plan considers:

- The top scenarios that drive risk to the organization
- Key questions that will need answers in those scenarios for operations, security, compliance, regulatory, fiduciary, communications, and other responsibilities
- Collection requirements and strategies to ensure proper root cause analysis and the answering of the key questions
- Roles and responsibilities of the individuals in the organization and partner organizations, such as incident responders, with specific action authorities identified
- The priority sites and crown jewels of the organization

¹² This is also an excellent place to start the requirements for what would be needed in the Collection Management Framework. See “Collection Management Frameworks—Looking Beyond Asset Inventories in Preparation for and Response to Cyber Threats,” www.dragos.com/wp-content/uploads/CMF_For_ICs.pdf

A key aspect of ICS Critical Control No. 1 is that it establishes a shared view on the risk and the outcomes the organization wants to achieve. As a result, it drives the requirements for the implementation of the other critical controls.

ICS Critical Control No. 1: Additional Learning and Insights

An individual power producer experienced an event where a gas turbine that was down for maintenance restarted without operator direction. Due largely to the growing complexity of industrial automation environments, the plant personnel were unable to get root cause analysis of why the generator started seemingly on its own. The personnel handled the situation professionally and engaged their ICS incident response firm. The incident response firm and the asset operator had already planned ahead of time on how to respond together. This planning allowed the incident response team to quickly deploy to the plant and upon analyzing the network traffic, determine there were commands coming from the human machine interface (HMI) that had activated the gas turbine control system. The teams determined that an upgraded HMI was now touch screen and the screen had been set to high brightness. After some expert sleuthing, the teams proved that a moth had entered the building at night, inadvertently hitting just the right part of the screen to activate the control loop. Safety and reliability were maintained, and root cause analysis was achieved in large part to proper planning and the capability to analyze ICS protocols and system-to-system analysis during the event. Figure 3 shows the ICS Incident Response Stages.



Figure 3. ICS Incident Response Stages

Control No. 2: Defensible Architecture

A defensible architecture is an architecture that reduces as much of the agreed-upon risk as possible through system design and implementation while also facilitating the efforts of human defenders. There is no such thing as a secure system or architecture—it is the human element that allows a defensible architecture to become a defended architecture. Many high-level frameworks and architectures can be employed, ranging from the Purdue Model to ISA/IEC 62443 architectures. The point is not the framework but the implementation of it to facilitate the security of the organization, taking into account the scenarios agreed upon with the first critical control. Organizations tend to design complex systems over time based on:

- System purpose
- Risk of impacts to system or threats to system
- Operational needs
- System communication needs
- User needs
- Vendor/manufacturer recommendations
- Regulations

It is the human element that allows a defensible architecture to become a defended architecture.

Considering the drivers and constraints that lead to an implementation, organizations often find that there are architecture choices appropriate for some sites but not all the sites around the organization. Having one standardized way to accomplish this architecture may not be possible or appropriate across all sites, but the architecture must facilitate, at a minimum, the data collection required from the scenarios identified in ICS Critical Control No. 1.

Common attributes of defensible architectures include:

- Asset identification and inventory for at least the crown jewels of the key sites
- Segmented environments where possible to reduce ingress and egress into as few pathways as possible, ultimately creating “choke points” for enhanced security and monitoring
- Determining when bi-directional access is needed, both now and in the future vs. truly read-only applications
 - For example, air gaps are not realistic in almost all environments outside of nuclear power plants. Modernization efforts and data access requirements significantly limit the ability of data diodes to be deployed across many sites. Data diodes can be successful in specific use cases, however, such as remote diagnostics monitoring of gas turbines where no other control or return access is required. In most organizations, a switched network and proper application of firewalls are common.

- Ability to collect network traffic and systems communication, such as managed network infrastructure with switched port analyzer (SPAN) ports or tap infrastructure
- Log collection from systems of value such as host-based log collection on HMIs and EWS, Sequence of Event logs from supervisory systems, and event and access logs from industrial equipment that supports it such as Syslog from PLCs
- Ability to go into a “defensible cyber position,” where enhanced connectivity and devices unnecessary for constrained operations are reduced during heightened situations, such as incidents in line with identified scenarios from ICS Critical Control No. 1

ICS Critical Control No. 2: Additional Learning and Insights

Review these two defense use cases related to the Ukraine 2015 and 2016 events: “ICS Defense Use Case 5: Analysis of the Cyber Attack on the Ukrainian Power Grid” and “ICS Defense Use Case 6: Modular ICS Malware.”¹³ These defense use cases highlight a need for improved architectural capabilities to enhance defender capabilities in reference to monitoring, alerting, mixed trust authentication and access capabilities, and the need for expanded system management supporting industrial DMZ architectures.

Control No. 3: ICS Network Visibility and Monitoring

The system of systems nature of an ICS drives a need for network monitoring to understand the interactions among those systems. ICS-specific monitoring includes the deep packet inspection of ICS protocols native to that environment. This security control carries numerous benefits, such as providing collection of data and detection of risk scenarios identified in ICS Critical Control No. 1, the ongoing validation of the architecture outlined in ICS Critical Control No. 2, and the placement, enhancement, and enforcement of ICS Critical Controls Nos. 4 and 5. Additionally, in a growing complex industrial automation environment, the ability to get root cause analysis even in non-cyber-related events is growing more difficult. ICS-specific network monitoring can aid in general resilience and recovery to avoid costly downtime scenarios and investigations.

ICS network visibility and monitoring is not just a technology problem. Among the five ICS Critical Controls, ICS Critical Control No. 3 is most often approached by organizations with the question, “what product do we buy to solve our problems?” There is no silver bullet technology that addresses this security control. An organization needs to consider the following factors to inform a technology selection:

- What data acquisition capabilities exist or are planned in connection with ICS Critical Control No. 2? (Consider endpoint/host acquisition, limited network collection, full network communications, multiple network visibility, ingest capabilities from other tools, and enriched analytics from additional providers.)

¹³ SANS “Industrial Control Systems Security” offers free resources about defense use cases. See www.sans.org/industrial-control-systems-security/

- What vendors and protocols are in use across systems of interest?
- What workforce staffing and capabilities exist or are anticipated to support the program?
- What processes exist or are anticipated in connection with ICS Critical Control No. 1 that will drive incident response actions?

With an understanding of internal organization capabilities and cyber maturity, an appropriate selection of technology can be pursued that aligns with the overall program goals shown in Figure 4.

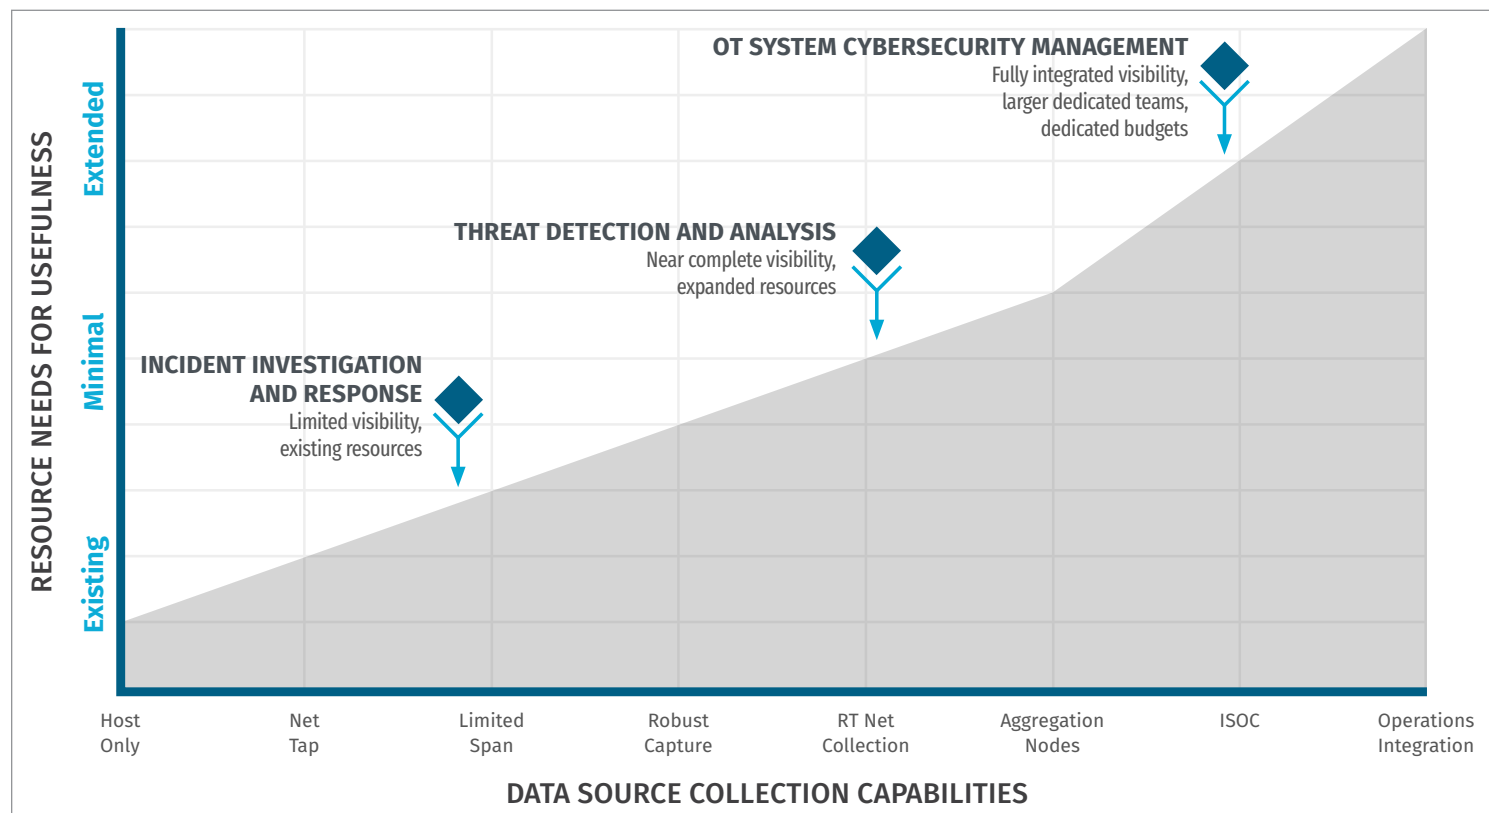


Figure 4. ICS Visibility, Maturity, and Capability Considerations

When evaluating technology platforms for ICS-specific network monitoring, look for the following features:

- Passive monitoring that is not intrusive to the industrial operations
- Asset inventory and topologies
- Identification and dissecting of ICS protocols to understand control communications
- Vulnerability identification
- Threat detection through key threat behaviors and tactics, techniques, and procedures of adversaries aligned with the risk scenarios to drive efficient and effective incident response
- Data collection and aggregation to support investigations and incident response
- Support for root cause analysis of operational issues and outages

Security teams are often concerned about the number of alerts they receive through security technologies such as those in the network monitoring category. Due to the industrywide lack of consensus about which security risks to monitor, teams look for any anomaly, which is a daunting and non-scalable approach. Instead, security operations personnel should focus on alerts that match the tactics, techniques, and procedures appropriate to the risk scenarios identified in ICS Critical Control No. 1. By doing so, security operations can create guides or playbooks against specific alerts, providing an efficient and repeatable approach to security monitoring.

ICS Critical Control No. 3: Additional Learning and Insights

Any attack demands reliance on all five ICS Critical Controls. They must be orchestrated in a manner that enables defenders and operators to act. As an example, the 2017 Saudi Arabian petrochemical facility attacks that targeted the process safety instrumented systems leveraged the TRISIS/TRITON malware. The custom attack capabilities delivered and modified within the operational target would have potentially been identified as the adversary communications, information collection, and ICS asset manipulation was occurring. This attacker activity would not necessarily trigger an existing malicious signature. It would more likely appear as an indicator of potentially suspicious communications due to configuration changes, baseline communications deviations, or other activity not in line with work management system orders approved by operations management. System defenders need the tools, technology, training, and operational processes to act on this demonstrated adversary approach.

Control No. 4: Secure Remote Access

The digitization of ICS and business requirements has demanded a significant increase in remote connectivity. Sometimes this remote connectivity is unnecessary and can be reduced or eliminated. In most industrial organizations though, remote connectivity is unavoidable and can have significant business and operations value.

Although the benefits of remote access are vast, so are the risks. Many critical infrastructure organizations moved to operating models in 2020 due to the global pandemic and the need to manage the human health safety risks. Prior to that time, such operating models would not have been allowed or accommodated at most organizations. Seemingly overnight, however, remote access has become the new normal. As a result, adversaries increasingly target the methods of remote access into industrial operations directly. It is no longer necessary in most companies to target the IT networks to get to the OT networks. Even when adversaries do target those networks, they may not be the organization's IT networks but instead the IT networks of their vendors, maintenance personnel, integrators, and equipment manufacturers. The adversary uses them to pivot directly into the OT networks. Establishing secure remote access is a must in modern-day industrial operations.

Multi-Factor Authentication (MFA) is a specific type of secure remote access security control that can be safely applied to most ICS environments. It has been shown to significantly reduce the number of adversary attack paths. Externally accessible connections should be the focus of the remote connectivity and application of MFA should be to externally accessible connections. Site-to-site communications that do not traverse the internet can be applicable but are not the priority. The priority is for remote connections that traverse shared private networks between two organizations, public networks, and the internet for access such as remote work by employees, integrators, original equipment manufacturers, and other vendors and partners. Wherever MFA is not possible, organizations should develop compensating controls designed with ICS Critical Control No. 2 in mind. Those compensating controls would include: jump hosts and opportunities for communications “break and inspect”; guiding remote connections through choke points for increased monitoring; and the capability to cut communications in heightened scenarios.

ICS Critical Control No. 4: Additional Learning and Insights

Testimony of Colonial Pipeline CEO Joseph Blount on June 8, 2021, provided details on the May 7 attack. He stated that the initial investigation identified adversary access through a legacy VPN path that was no longer being used. While efforts to update and move to stronger remote access capabilities had already been pursued and implemented, the previous system had not been decommissioned. This is a lesson worth highlighting because the challenge of remote access is to first understand where connectivity exists and manage efforts to improve security controls. This challenge is one of the reasons for the ordered nature of the five Critical Controls and why secure remote access is focused on after visibility and monitoring. An organization must also ensure that projects include removal of previous technologies at appropriate locations. Frequently, as new systems or programs are implemented and during cut-over testing periods, both systems remain active due to the critical nature of the operational environments, leaving older systems not appropriately decommissioned. In addition, there may be shadow remote access throughout operational networks located at field sites with weaker remote access capabilities for support teams to directly access. There may also be remote trusted connections to a third-party vendor environment with a weaker remote access program that could in turn act as a path into the organization. A complete inventory of remote access paths is absolutely required to ensure all entry points are protected.

Control No. 5: Risk-based Vulnerability Management Program

A risk-based vulnerability management program focuses on those vulnerabilities that actually drive risk to the organization, especially those that map to the scenarios identified in ICS Critical Control No. 1. Often, the vulnerabilities that drive risk in ICS are those that help an adversary gain access to the ICS or introduce new functionality that can be leveraged to cause operational issues such as the loss of view, control, or safety. The focus of the vulnerability management program is not simply to patch vulnerabilities but also, in many cases, to mitigate their impact or monitor for their exploitation. Each year only roughly 4% of the vulnerabilities in ICS environments are required to be immediately acted upon based on a risk-based approach of what adds new functionality to the ICS or is already under active exploitation.¹⁴ Upwards of 10% of identified vulnerabilities each year are completely useless in that they cannot be weaponized effectively or are just hype and incorrect in their advisories. The remaining vulnerabilities can either simply be monitored for abuse or mitigated entirely through simple actions such as the disabling of services and communication paths that are unnecessary on the device or at the boundary device such as a firewall through application of ICS Critical Control No. 2.

Often the focus on vulnerabilities drives conflicts between IT and OT staff because finding and patching every vulnerability in an operational environment across equipment with deployments of upwards of 30-year life cycles can be overwhelming. This is especially true when patch application can have unknown effects or require reboots or maintenance periods that may not regularly happen at industrial assets. Focusing instead on the key vulnerabilities with a focus on a risk-based approach with the application of ICS Critical Control Nos. 2 and 3 allows for the tension, workload, and potential for disruption to be significantly reduced.

Some regulatory approaches such as NERC CIP provide requirements around security patch management and the corresponding change management criteria required whenever security-related patches and security control changes could be affected. While these requirements are absolutely necessary and provide a security benefit to entities subject to the regulation, they come with a lot of pain with regard to the time frames in which they need to be assessed, performed, and documented on an ongoing forever basis. In addition, there is no real ability to assess the priority or risk of identified security patches because all identified applicable security patches are treated equally and require action, which may put the systems at risk to implement. The high frequency of occurrence

¹⁴ The Dragos Year in Review reports provide analysis on the threats and vulnerabilities in the community. Across each year, the report finds that about 4% of the vulnerabilities add new functionality to the environment that an adversary can abuse or has already exploited. A significant portion of vulnerabilities in ICS simply do not reduce the risk to patch or mitigate and instead should generally just be monitored for exploitation. To read the 2021 report, visit www.dragos.com/year-in-review/?utm_campaign=Q121%20-%202020%20Year%20In%20Review&utm_source=SANS%20ICS%20Summit%20Keynote

requirements for patching or vulnerability management programs forces entities to interact and modify operational systems, introducing additional system risk. For these reasons, entities need a way to examine the applicable security patches and make an operational risk decision about when and how to address the identified vulnerabilities. An example of this approach can be seen in Figure 5.

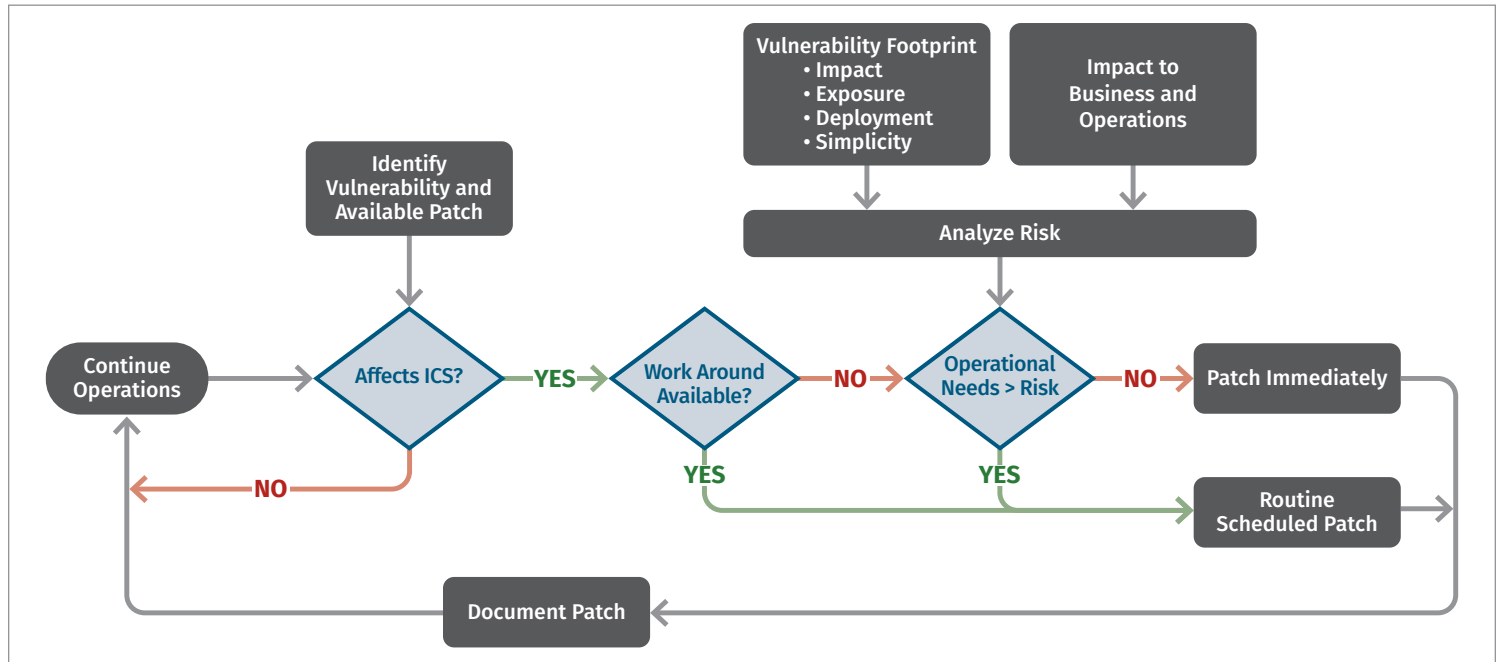


Figure 5. DHS Patch Urgency Decision Tree¹⁵

Risk-Based Vulnerability Management Program Tips

- Key vulnerabilities are those that support adversary operations against the scenarios prioritized by the organization, add new functionality to the environment that the adversary could reasonably utilize, or are actively being exploited. If any of these three qualifications are met, it is likely a key vulnerability that should be mitigated or monitored.
- Where possible, utilize software bill of materials (SBOM) to identify the underlying vulnerability and mitigate it, regardless of whether other vendors have issued an advisory.
- It is incredibly common in the ICS community for multiple vendors to have the exact same vulnerability yet only one vendor issues an advisory and mitigation is based on the reporting researcher contacting them. Such was the case in the PIPEDREAM malware that took advantage of a specific version of Codesys software embedded in hundreds of different programmable logic controllers (PLCs), but only a few vendors disclosed the problem.

¹⁴ www.cisa.gov/uscert/sites/default/files/recommended_practices/RP_Patch_Management_S508C.pdf

- Sometimes patching will not reduce the risk. Worse, it may introduce additional risk. Mitigating the vulnerability may simply require changing a firewall or monitoring for exploitation.
- Actively querying devices for information to support vulnerability management can be disruptive or even dangerous; most modern technologies in ICS Critical Control No. 2 can passively identify vulnerabilities to support the program. If utilizing active querying, ensure that testing is done first and attempt to confine the querying window to maintenance windows or downtime at the facility. It's always a best practice to avoid seeking information through active querying that you cannot utilize anyway. As an example, if knowing the firmware version of the card on the controller is important to determining whether it is vulnerable but no action is going to be taken on the controller, simply apply the mitigation (such as changing a firewall or monitoring for exploitation) and do not obsess about identifying the firmware version. No need to add more risk to the operations you are attempting to mitigate.

ICS Critical Control No. 5: Additional Learning and Insights

Although many organizations release and track ICS device- and application-specific vulnerabilities, the Dragos Year in Review report provides analysis of those vulnerabilities. In the 2021 report, the Dragos team provided the analysis shown in Figure 6.




Figure 6. DRAGOS Year in Review Vulnerability Analysis Findings

For 77% of the vulnerabilities examined, the report indicated that the vulnerability required direct access to the devices within the ICS segment. For many of these vulnerabilities, an adversary could achieve the same objective with direct network access to the vulnerable device by simply crafting commands, manipulating traffic, or engaging in other attack approaches without exploiting the vulnerability. Said another way, if the vulnerability was patched, could the same outcome be achieved through another means with direct ICS network access? If the answer is yes, identify which ICS Critical Controls would mitigate the attacker effect (No. 2, No. 3, or No. 4). On the other hand, if the vulnerability is a unique capability and exploitable, then patching or mitigation strategies should be pursued if the risk of exploit is greater than the risk of operational impact due to the patching or mitigation approach.

Summary

The five ICS Cybersecurity Critical Controls discussed in this paper provide a path for critical infrastructure organizations to pursue through new capital investment projects and in programmatic operational and maintenance initiatives. They can be pursued in order and in concert with one another to create a robust ICS cybersecurity program that is tailored to the risks facing the organizations. These prioritized critical controls can help guide organizations seeking recommendations and guidance on what to do next based on threat-informed activities instead of over- or under-investing. Critical Controls supporting elements need to include those in Figure 7.

While the five controls highlighted in this paper will act as a valuable resource to practitioners and leaders alike, they are less effective and more difficult to implement without the appropriate support and culture within an organization, including:

- Identification of critical facilities
- Operations-aligned response plans
- Organization coordination on risk-based scenarios

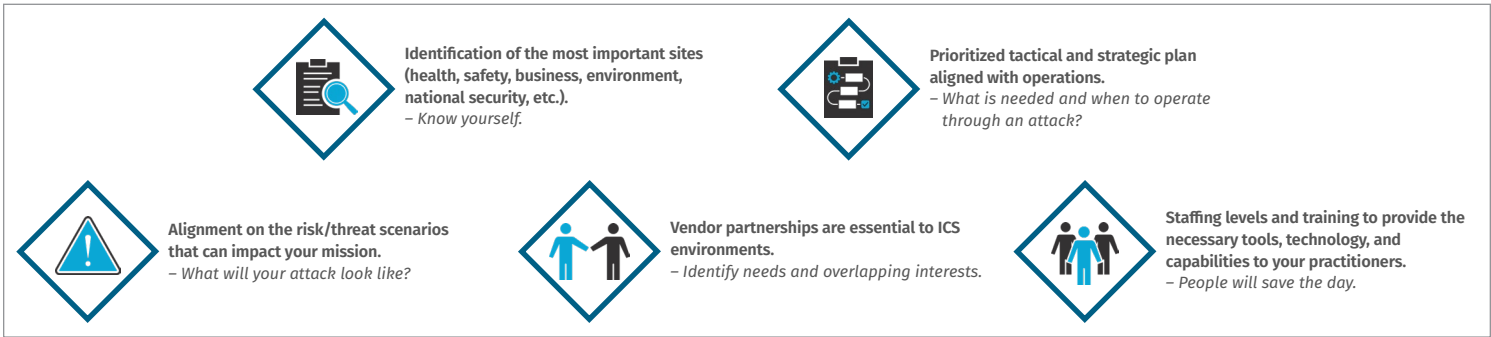


Figure 7. Foundational Support Elements for ICS Critical Controls