Industrial Control Systems Security Checklist

A guiding document for developing secure Industrial Control Systems.

Written by: Lars Flå, Eric Törn, Knut Vidar Skjersli and Fredrik Bakkevig Haugli

January 2023

V1.0

SINTEF

Foreword

With a more connected world, where someone can attack digital industries in distant countries with low risk and cost, cyber security is becoming increasingly important on a nation-wide level. Several recent attacks illustrate this, for example towards the power grid in Ukraine in 2015 and 2016, a petrochemical plant in Saudi Arabia in 2017, the attack towards a cloud service provider to the train system in Denmark in 2022, and others. These attacks can, if successful, harm important industry and infrastructure such as hospitals, transportation, and industrial manufacturing.

While traditional Industrial Control Systems (ICS) have been mostly isolated, they are increasingly getting connected to the internet, increasing the risks of cyberattacks.

This guidance document intends to contribute to securing ICS and reduce the probability of successful cyberattacks and limit consequences affecting health, environment, and the economy.

This guidance document is a deliverable from the project *Ragnarok* at SINTEF, and complements the *Security Aspects of Industrial Control Systems* document and the IoT checklist document, also a deliverable in this project. In this project we looked at both ICS and IoT security, both at the state of the art and within the context of industry 4.0 and the potential collapse of the SCADA pyramid.

Changelog

v1.0

Audience

This guidance document is primarily aimed at ICS integrators and asset owners.

Scope and ambition

The goal of this document is to provide a checklist for ICS security that ICS integrators and owners can use when developing and operating their systems. The document draws inspiration from guidelines and standards on ICS security, but the points raised represent the view of the authors. We acknowledge that this is by no means an exhaustive list and believe that such a list would be difficult to create. ICS is used in a wide variety of industries, each facing potentially different situations and security needs, involving a mixture of people, processes, and technology.

Instead, we highlight some general concerns and present it as a starting point for securing ICS. This guidance document is in many ways not sufficiently detailed to be used as a sole source for ICS security, and users are encouraged to consult additional sources on the aspects of interest.

Security Levels

We define three security levels, partly inspired by the OWASP Application Security Verification Standard and partly inspired by the security levels in IE 62443. The security levels map to what we generally believe should be implemented for all ICS use cases (level 1), what should be considered on a case-by-case basis (level 2), and what should be considered for particularly critical use cases (level 3).

Level 1: this is the basic level of security that should be implemented everywhere. It ensures a security baseline which aims to defend against threat actors whose capabilities and motivation limits their scope to easy-to-find and easy-to-exploit vulnerabilities. One should however also assess the need for mechanisms on level 2 and 3.

Level 2: These are the mechanisms which one should consider implementing, based on case specific criticality and threats. They take more advanced threat actors into account, who may have longer time horizon when planning and executing attacks, and target vulnerabilities that are harder to find and harder to exploit.

Level 3: These are the mechanisms which generally are intended for particularly critical use cases. They consider that a threat actor may have significant to unlimited resources and motivations, typically threat actors associated with a nation state.

Overall principles

Q1. Is ICS security guided by a risk-based approach?

ICS security efforts should be guided by a security risk management process. The work on security risk should be a continuous process involving monitoring of risk, assessment of risk and response to risk. These steps should be supported by an activity defining priorities and constraints for risk decisions (e.g., in case of conflicting requirements, which should take precedence). As security and safety risks may have reciprocal implications, the two processes should inform each other.

Q2. Is there a security program in place?

A company should develop a security program defining the policies, procedures, and activities involved in handling the security of an ICS. A security program should be anchored in top level management and the security effort should include cross functional teams (including representatives from IT, process control and company risk management). The company should identify potential synergies with existing security programs for IT and explore the potential for sharing resources.

Q3. Do you have procedures in place to regularly take backups of ICS configurations?

Such backups may prove invaluable when recovering from a security incident, especially if the ICS is hit by crypto malware. Further consider automating the backup and backup verification processes.

Q4. Do you have an overview of the device and software assets in the ICS?

An inventory of all devices and software in the ICS should be created and updated. If automated solutions are used, care must be taken to ensure that they do not interfere with normal operations.

Q5. Is the ICS designed with defense-in-depth in mind?

ICS security should utilize a defense in depth concept where one uses multiple layers of security, such that if one layer fails, another layer is still securing the system. Such layers can be of a technical, process, or human nature. As an example, we consider different solutions for preventing unauthorized access to an ICS network: A procedure requires that all remote connections should pass through the security mechanisms of the IT network (before reaching the ICS network). An intrusion detection system on the IT network monitors traffic for suspicious remote connections. An on-site human is required to approve any remote connection requests.

Network

Q6. Is the network properly segmented?

An ICS should be partitioned into multiple sub-networks or zones based on security requirement. The number of sub-networks or zones depends on the complexity and security requirements of the ICS. As a minimum, control networks should be separate from non-control networks and safety instrumented systems should be located on a separate network. A further improvement is to segment the ICS into more zones based on risk analyses for the ICS. There should be a firewall between different sub-networks.

Q7. Is the network properly segregated and configured according to the principle of least privilege?

The network should enforce rules for communication between services and hosts. Only the devices, users, and services with an operational need to communicate with each other should be allowed to do so. Ports and routes should be enabled only as it becomes needed, and one should consider implementing rules regarding the protocol, type of data, or amount of data a system can send.

This can be implemented with firewall policies.

Q8. Are your firewalls using white-listing rather than black-listing?

All communication through firewalls should be blocked by default and exceptions should be made only when needed.

Q9. Are you periodically testing your firewall policies?

Firewall rules should be tested periodically to reveal any deviations from the defined set of rules. Such deviations can easier be detected if current firewall configurations can be compared to backups taken before the firewall was commissioned.

Examples of tests are vulnerability scanning and penetration test, which can be automated to run at certain intervals.

Q10. Have you limited access to internet for devices on the control network?

Internet access by devices on the control network should be strongly discouraged, and any traffic out of the ICS should go through the DMZ, see Q7. Direct internet access for devices on the control network may constitute a considerable security risk and undermine the defense in depth approach.

Q11. Are safety instrumented systems separated from the rest of the ICS?

Safety Instrumented functions should be deployed on a separate network and not be reachable from the rest of the ICS. This reduces the risk of a cyberattack causing physical destruction, as it would require the attacker to control both the process control system and the safety instrumented system.

Q12. Are you using one or several DMZs?

Traffic should be prevented from transiting directly between control and IT networks. Instead, all traffic should terminate in a DMZ. Deviations from this should only be granted in very specific cases. There should be documented justification with risk analysis, business benefits for doing so, and a person responsible for each permitted incoming and outgoing data flow. Normally, all services and components that are meant to be accessed both from the ICS network and the IT network should be placed in this zone. Examples include remote access solutions, patch management servers and historians. Outbound packets from the control network or DMZ should be allowed only if those packets have a correct source IP address that is assigned to the control network or DMZ devices.

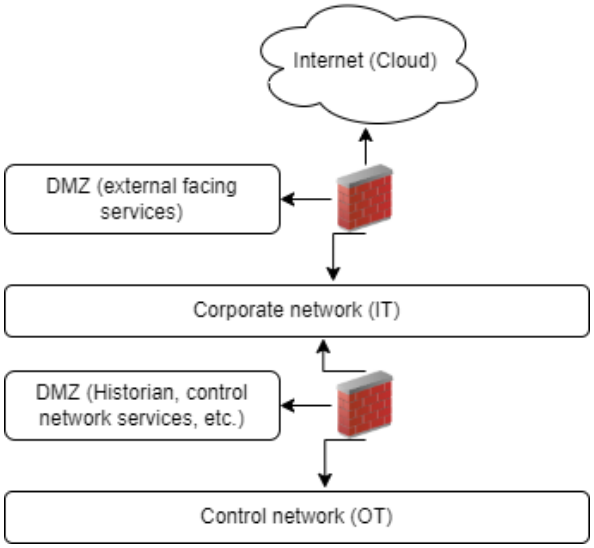

Figure 1: High level security architecture

Q13. Are you using packet-filtering firewalls?

Consider what type of firewall functionality is needed for your ICS, but rules should as a minimum be both IP address and TCP/UDP port specific. Filtering packets on the network and transport layers based on IP addresses and ports is the simplest type of firewall. More advanced firewalls also include filtering based on the application layer, see Q32.

Q14. Have you created a firewall policy and is this kept up to date throughout the life of the firewalls?

A policy outlining how firewalls should handle traffic between different parts of the ICS and between ICS and the DMZ/IT should be created. The policy should be kept up to date to accommodate new hosts and services being introduced in the ICS and to serve as a reference during inspections of the firewall rules that are in place at any given time.

Q15. Are outbound rulesets receiving as much attention as inbound rulesets?

The firewall rules regulating outbound communication should be configured with the same rigor as rules for inbound communication. The reason for this is that attackers may want to export data or that the attack relies on outbound communication for command and control. All outbound traffic from the control network to the IT network should be source and destination-restricted by service and port.

Q16. Are you using a proxy server for incoming connections from the internet?

Request for information or services from the ICS domain (e.g., files or connections) should be handled by proxy servers placed in DMZs or other zones with suitable security requirements for this purpose. Processing external requests in this way manages complexity and provides additional protection.

Q17. Are you preventing ICS components from being discoverable on the network?

Configure router and switches in the network to drop any packets that can be used for network discovery (e.g., ping and arp-ping).

Q18. Do protocols in use have security features for encryption, authentication, access control and logging?

Some of the protocols used in the ICS domain do not implement security. Investigate whether the protocols in use support security features, such as authentication, access control and logging, and whether such features can be enabled without affecting operations. As an example, verify whether authentication introduces unacceptable levels of latency to the communication. When implementing security, one should consider threats such as information disclosure, tampering, replay, spoofing, and denial of service (for instance through message flooding). Examples of security for different protocols include Transport Layer Security (TLS) for TCP, built in security features in OPC UA and security plugins for DDS.

Q19. Can firewall configurations only be changed by users with special privileges to do so?

All firewall management traffic sent over the regular network should be protected (using standard communication security such as encryption and authentication) and requiring the user to authenticate to get access to the functionality. There should be a limited number of users with these privileges. Traffic should also be restricted to specific management stations.

Q20. Do you protect against ARP based man in the middle attacks?

The ICS should protect against ARP based man in the middle attacks launched by an attacker who has obtained a presence on the local network. This attack can lead to both denial of service, tampering of communication and information disclosure.

ARP based man in the middle attacks can be prevented by using static ARP tables, message authentication or by monitoring for ARP poisoning.

Q21. Is your firewall management traffic carried on a separate, secured, management network?

All firewall management traffic should be carried on a separate, secured management network (e.g., out of band[1]) requiring authentication. Traffic should also be restricted to specific management stations.

Q22. Are you able to isolate networks in case of an attack?

There should exist a functionality to isolate networks from other networks in case of an attack. The interfaces between the IT and ICS networks, and the interfaces between the safety instrumented system network and the remaining ICS network are prime candidates for such functionality. Isolating networks can help ensure the continued safe operation of the ICS network, but care must be taken to ensure that the action itself does not interfere with operation.

Q23. Are you using different application-level protocols between the DMZ and the corporate and control network?

Protocols in use between the control network and DMZ should not be allowed between the DMZ and IT network. The same applies in the opposite direction. This practice improves security by requiring an attacker to exploit vulnerabilities or functionality in two different protocols to be able to pivot from the IT to the ICS network.

[1] https://en.wikipedia.org/wiki/Out-of-band_management

Q24. Are your switches using MAC address locking?

MAC address locking prevents attackers from connecting unauthorized devices to the network. It also ensures that the network topology does not change without the network administrator's knowledge (e.g., legitimate changes to the ICS).

Q25. Have you configured boundary protection devices to fail in a predetermined state?

Failing boundary devices (for instance due to hardware failures) should fail to a predefined state. If operational concerns allow, boundary devices should fail to a state where all traffic is denied. This may however not be possible due to for instance safety concerns.

Q26. Have you implemented unidirectional gateways where there are one-way data flows between security domains?

Unidirectional gateways (or data diodes) only permit data to flow in one direction, often by restricting flow in the reverse direction on the physical layer. However, to enable reliable transfer of data, many implementations allow acknowledgement to be transmitted in the reverse direction. By only allowing acknowledgements to flow in the reverse direction, the attack surface is reduced. Consider if this may be applicable to some of the zones in your ICS.

Authentication and authorization of users

Q27. Do you have an authentication and authorization system for assigning users only the necessary levels of privilege?

Users and user roles should only have the rights needed to perform their task; a practice known as least privilege. There should exist procedures to efficiently check what users have access to what systems, change permissions for existing users, add new users, and remove users. This includes removing inactive users and users with default credentials (for instance left behind by the vendor or system integrator). Users should by default not be able to access critical devices or systems on the control network. When such devices or systems must be configured, specific users should only be given access for a limited time.

Q28. Is two factor authentication required for all remote connections?

Any remote connection to the ICS should at least require two factor authentication. However, stopping at two factor authentication should only be the requirement if accessible functionality is not critical. Furthermore, any remote connection should first authenticate to the IT network, and then reauthenticate to access to ICS network.

Q29. Are remote sessions terminated after inactivity?

Remote sessions should require users to reauthenticate after a certain time of inactivity.

Q30. Is two factor authentication complemented with work permit and dual approval?

In cases where remote users have access to critical equipment, for instance safety instrumented systems, access should only be granted if two factor authentication is complemented by issued work permits and dual approval. Work permits aim to ensure that access to the ICS network to make changes is only permitted if there is a permit documenting permission to do so. Dual approval ensures that all changes must be approved by two people.

Monitoring and auditing

Q31. Are you monitoring the ICS network for anomalous communication?

Consider deploying Intrusion Detection Systems to detect anomalous behavior. Intrusion detection systems can typically be rule based or anomaly based. Rule based IDS generally suffer from only being able to detect known attacks while anomaly-based IDS generally suffer from false positives.

Q32. Are you using firewalls with deep packet inspection on the application level?

Consider if filtering and stateful firewall should be enhanced with deep packet inspection on the application layer. On this layer, packets can be filtered based on protocol and application type and specifications. This can identify vulnerabilities and suspicious activity that cannot be detected by devices operating at the network or transport layers. The limited number of formats, especially the prohibition of free form text in email, eases the use of such techniques at ICS boundaries. Filtering on all levels (network, transport, application) offers good security, but can increase latency, configuration complexity, and cost.

Q33. Can you backtrace and analyze attacks?

The ICS should be able to log and timestamp key events, including access, operating system events and process configuration changes. This includes if the actions were taken by a registered user. This is used both for auditing and backtracing purposes.

Software

Q34. Have you verified the source of your software?

You should establish procedures for verifying the integrity of software and software updates to protect the ICS against supply chain attacks. This can include having the vendor sign its software or by calculating the hash (e.g., SHA) of the software. Supply chain attacks may introduce backdoors into legitimate software to allow attacker to gain a foothold on the ICS network or change the software, so it interacts with the process under control in malicious ways.

Q35. Do you have a strategy for patch management?

You should establish a patch management strategy for the software used inside the ICS. Patching software in an ICS in operation may include shutdown of operations, thorough testing, and creation of backups. A risk analysis should guide the decision of whether and how quickly to patch. As it may be unfeasible to patch the system, for instance due to legacy products no longer being supported by the vendor, suitable compensating measures must be taken.

Q36. Are you able to detect and prevent malware?

The ICS should implement the functionality to detect and prevent malware. This can be achieved through, for instance, application whitelisting or through monitoring devices for new binaries or changes to existing binaries. Note that application whitelisting makes the process of patching more complex.

Hardware and devices

Q37. Are you limiting access to hardware only to people that require such access?

There should be mechanisms in place to ensure that only authorized personnel have access to the different components in the ICS. This can include measures to restrict physical access to rooms and cabinets containing equipment.

Q38: Have you disabled unnecessary functions, debugging access ports, TCP/IP ports, and services on your devices?

The exposure of unnecessary services by devices in the ICS network can constitute a significant cyber security risk. Disable or remove all services and protocols not necessary for safety or normal operations, including troubleshooting or debug services. Such services and

protocols can have been enabled by default from the device manufacturer, left by the service integrator after the commissioning of the plant, or left by the operator after regular maintenance.

Q39. Do you have user authentication for devices such as PLCs?

If possible, there should be an authentication mechanism in place for accessing and configuring devices on the control network. Most PLCs have an authentication system with username and password.

Q40. Are you able to determine if physical equipment has been physically tampered with?

This can be implemented in a physical way by placing seals and tape on equipment, or in an electronic way, by collecting and storing logs. Location is, however, important. Equipment placed far outside the control of the operator, for instance in a remote electrical substation, have different requirements than those residing in continuously manned control rooms.