National Security Agency/
Central Security Service

INFORMATION ASSURANCE DIRECTORATE

Securing Assets within a "Closed" Industrial Control System (ICS) Network Perimeter

This publication is the second in a series intended to help Industrial Control System (ICS) owners and operators in need of improving the security posture of their systems. This document will focus the reader on aspects of system security within a “closed” ICS network perimeter, and give them a systematic approach for implementing the access control concept of Least Privilege. By restricting the accesses for process automation devices to and from only authorized subjects (i.e. individuals and/or other devices), ICS security can be built in a layered manner. Starting at the heart of a closed loop system and controlling access to devices in the system network out to devices at the physically-controlled system network boundary, ICS owner’s can build hardened “closed” ICSs which both logically and physically address the perceived security threats.

Control Systems Division

Contents

1. Introduction 1

2. Definition of a Closed Network 1

3. Security Recommendations 2

3.1. Network Segmentation and Traffic Filtering 2
3.2. Authentication 6
3.2.1. User Authentication 7
3.2.2. Peer-To-Peer Authentication Services 8
3.3. Cryptographic Security and Data Confidentiality 9
3.4. Device Hardening 10
3.5. Security-Related Firmware and Software Updates 11
3.6. Security Event Monitoring 13

4. Preserving ICS Performance and Functionality 13

4.1. Bandwidth Limitations 13
4.2. Latency Requirements 14
4.3. Availability Requirements 14

5. Conclusions 15

1. Introduction

Industrial Control Systems (ICSs) control and monitor complex industrial processes like petroleum refinement, chemical production, product manufacturing, and electric power generation and transmission. Much of the United States' critical infrastructure is dependent on industries that employ networked ICS systems. Sabotage or disruption of these industries can have wide-ranging negative effects including loss of life, economic damage, property destruction, or environmental pollution. Our reliance on ICS networks makes them attractive targets for electronic attack.

In this publication, we discuss security recommendations appropriate for application in ICS network environments. The recommended techniques and practices provide sufficient security for closed ICS networks that are secured against physical or electronic access by unauthorized outsiders.

2. Definition of a Closed Network

A closed network is any network that is "significantly" protected from physical and electronic access by unauthorized individuals (outsiders). The following qualities minimize the likelihood of unauthorized physical or electronic intrusion by outsiders, and thus form the definition of a closed network:

- Assets on the closed network operate within an uninterrupted, defended physical perimeter (locked or monitored building complex, unbroken fence line with locked gates, etc.)
- Assets on the closed network are electronically connected to one another via continuous or intermittent, wired communications links carried on media located entirely within the secure physical perimeter (networked wireless devices do not fit this definition)
- All users that can access networked assets on the closed network have known, definable trust and privilege profiles

Many ICS networks violate the strict rules listed above. For example, ICS network segments may be distributed across multiple, geographically-separated sites that are connected by electronic wide area network communication links that exit the secure physical perimeter. Other common violations include the use of wireless network links within an otherwise closed network segment, or connection of ICS network segments to non-ICS network segments that directly or indirectly connect to the Internet (e.g.; network connections between critical SCADA assets and the corporate network segment for billing, trending, etc.). These links represent potential electronic entry points through which unauthorized outsiders may gain entry into the secured network perimeter. The security principles recommended in this document can be applied to such networks after these electronic entry points are sufficiently secured and monitored to significantly reduce the chance of unauthorized outside intrusion. We discuss recommended practices for securing the electronic access points that violate the rules defining a closed network in our publication "*Defending Risky Electronic Access Points into 'Closed' Industrial Control System (ICS) Network Perimeters*".

3. Security Recommendations

Many of the security recommendations discussed in this document are selected to enforce the principle of least privilege: users and devices on the network are only given access to the minimum capabilities required for each asset to satisfy their intended function. The goal is to limit access to available functions, actions, and sensitive data to those devices or users within the network perimeter that either have privilege to access the action/data or are trusted not to access the action/data. Central to this concept is the ability to profile all users and devices on the network and to identify the rights and privileges that must be granted to them. Once these requirements are defined, the techniques and practices discussed below can be applied to make it sufficiently difficult for users and networked devices to execute actions beyond their authorized capabilities. In addition, the techniques discussed below will also help secure the system against unauthorized outside intrusion (if the system becomes connected, i.e.; not "closed, to other systems/networks), and slow or eliminate the propagation of electronic attacks by motivated attackers or malicious software (e.g.; worms or viruses).

3.1. Network Segmentation and Traffic Filtering

Electronic attackers often utilize unsecured communications services to perform malicious actions. Blocking access to unnecessary communications services or functions whenever possible minimizes the avenues through which these electronic attacks can propagate.

Users and devices on the ICS network with similar roles, privileges, criticality, and communications requirements should be electronically partitioned into distinct sub-networks in order to localize the communications traffic associated with their primary functions. Ideally, two distinct sub-networks should only be connected to one another by a single network link through which all data passed between the two dissimilar sub-networks flows. These concentrated connection points provide convenient locations to apply traffic filtering technologies that restrict access to unneeded communications services and functionality. All additional connections between two distinct sub-networks (e.g.; for redundancy/reliability) must be equally secured.

We can identify communications classes common to ICS networks and use these to suggest network partitioning architectures that provide improved isolation of potentially vulnerable communications services.

- **Supervisory Control and Data Acquisition (SCADA):** plant operations personnel monitor the status of the plant's processes on PC-based operator workstations or system mockups. SCADA master terminal units (MTUs) periodically gather process status (e.g.; switch positions or motor RPM) from networked ICS equipment and make it available to operators and other personnel. In addition, the MTUs relay process control commands from the operations personnel to the networked ICS equipment to communicate requested process changes. ICS SCADA operators have authorization to view the process data gathered by the SCADA MTU and, usually, to issue commands to alter the state of the process. The transfer of process status and user commands between operator workstations, the MTU, and the process automation devices take place using protocols that often lack inherent security features like password protection or cryptographic authentication. SCADA, process automation, and database access protocols like DNP3, Modbus,

OPC, or SQL must be isolated to minimize the risk of unauthorized manipulation of process data or injection of process control commands onto the network.

- **Read-Only Access to Process Data:** some important ICS functions require read-only access to the process status data, but do not require authorization to issue process altering commands. Some common examples include process data archiving, trend analysis, contingency analysis, and billing. Read-only process data users should be blocked from accessing any process control commands available to more privileged SCADA operations assets (e.g.; the MTU or operator workstations).

- **Process Configuration/Engineering:** engineering access services provide the ability to change settings in critical process automation devices, often including the ability to alter process control logic and issue process control commands. These services are often protected by relatively weak user authentication mechanisms and very often transmit authentication parameters (e.g.; passwords) unprotected. Vulnerable engineering access services should be isolated to help prevent unauthorized manipulation of process control configurations and the potential damage and injury that such actions can cause. Some ICS device vendors implement engineering access services by mapping device settings to a SCADA protocol. Often, such interfaces can serve as either a traditional SCADA interface (see above) serving process status and process control commands to/from the control center, or as an engineering access interface offering access to critical device settings. Such services should *always* be considered engineering access interfaces, and isolated accordingly.

- **Closed Loop Process Automation:** industrial processes are controlled by process automation devices[1] that gather process status, run preconfigured process control algorithms, and send commands to implement the process changes dictated by the executed control algorithms (e.g.; raise motor RPM, open a breaker, or turn a pump on). Cyclic, closed loop control cycles may be on the order of tens of milliseconds, so all communications associated with this process must be high performance, low latency, and high reliability. Closed loop process automation communications between process automation devices and intelligent, networked process actuators (e.g.; pumps, valves, electrical breakers)[2] should be isolated to help preserve communications performance. Preventing remote electronic access (e.g.; via the control center SCADA network) to process actuators allows system designers to configure process automation devices to block execution of command sequences that would result in dangerous or unsafe conditions. For example, a command received over a SCADA link to turn on a pump can be supervised by logic in the automation device that received it to prevent the pump from operating if the tank is already full. Ensuring that the pump is not directly addressable via the SCADA link prevents an attacker from bypassing this safety logic.

[1] Process automation devices include Programmable Logic Controllers (PLCs), Remote Terminal Units (RTUs), and digital protective relays that monitor and protect electric power system equipment (transmission lines, transformers, etc.)

[2] Many modern process sensors and actuators, like fluid level meters, pumps, electrical breakers, and valves, have digital communications interfaces rather than simple analog control or measurement wire connections. Many of these process automation bus protocols can be carried on TCP/IP LAN segments.

Figure 1 shows an example ICS network architecture that follows the network isolation suggestions discussed above. In this example network, all SCADA and engineering access communications are implemented using TCP/IP communications services. TCP/IP SCADA solutions have become much more common in recent years and, in many ways, are overtaking serial implementations in popularity. There are many SCADA protocols designed to communicate over TCP/IP network infrastructures, including DNP-IP, Modbus TCP, IEC 60870-5-104, and IEC 61850. Similarly, ICS device vendors are routinely including TCP/IP engineering access services in their more recent products.

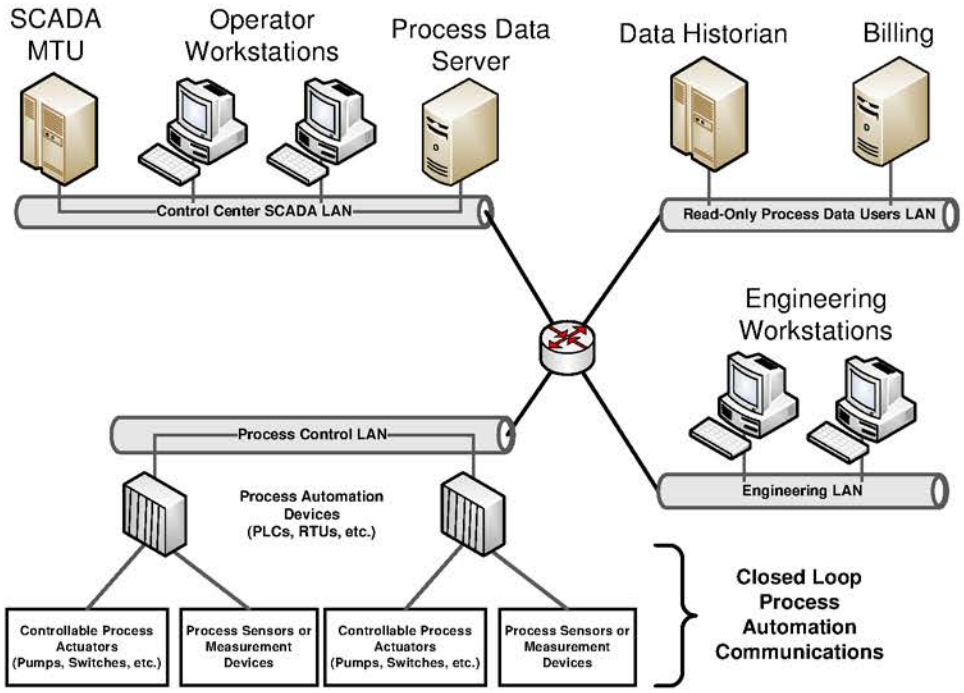


Figure 1 Segmented ICS Network with TCP/IP SCADA and Engineering Access Links

In the proposed network architecture shown in Figure 1, separate TCP/IP local area network (LAN) segments have been created to isolate the Control Center SCADA assets (e.g.; SCADA MTU), the Process Control assets (e.g.; RTUs), Read-Only Process Data Users assets (e.g.; Historian), and Engineering assets (e.g.; device configuration/engineering workstations). In addition, any intelligent, networkable process sensors and actuators available within the closed loop process automation logic have been attached directly to process bus communications interfaces on the process automation devices, and are not directly addressable via the Process Control LAN.

In Figure 1, a four-interface TCP/IP router (or firewall) should be placed between the four LAN segments. Access control lists (ACLs) programmed into the router by a network security administrator are used to strictly enforce which types of TCP/IP traffic are allowed to flow between the four LAN segments. Custom ACLs should be written to block all network traffic except that required to conduct the normal operation of the ICS network. The necessary ACL entries depend greatly on the requirements of the implemented network, but the following table provides a high level summary of the suggested outcome of an effective set of ACL rules for the segregated network in Figure 1.

Suggested Router Access Control List Rules	Comments
ICS device engineering services should only be allowed to flow between the 'Engineering LAN' and the 'Process Control LAN'	All process automation device configuration activity should take place on the 'engineering LAN', allowing these potentially vulnerable services to be blocked from originating from any other LAN segment. Any service offering access to critical ICS device settings or configurations should be considered an engineering access service and isolated according to this suggestion.
Process data access services flowing between the 'Read-Only Process Data Users LAN' and the 'Control Center SCADA LAN' should not allow changes to process data values or execution of process commands	It is prudent to serve the read-only process data from a server on the 'Control Center SCADA LAN' that is physically separate from the SCADA MTU. It is also easier to choose a process data server that can run secure data access protocols than it is to expect the SCADA MTU to offer such services. This process data server should maintain a current mirror of the process data values, periodically updated from the SCADA MTU. The process image in the process data server can then be offered to read-only process data users via a secure process data transfer service (e.g.; via a secured database service). Allow only secure database traffic to flow directly between the 'Control Center SCADA LAN' and the 'Read-Only Process Data Users LAN'.
SCADA services should only be allowed to flow between the SCADA MTU on the 'Control Center SCADA LAN' and the process automation devices on the 'Process Control LAN'	SCADA services almost never include inherent security features like passwords or cryptographic authentication. Access to such services, especially those interfaces that offer access to process control commands, should be blocked whenever possible. SCADA protocols may be used within the 'Control Center SCADA LAN' (e.g.; to mirror a subset of the process data image to the process data server) but only the SCADA MTU on the 'Control Center SCADA LAN' should be allowed to access the SCADA services offered by process automation devices on the 'Process Control LAN'.
All non-essential services should be blocked from flowing between network segments	Blocking non-essential services helps prevent attackers from exploiting insecure services to propagate attacks from one LAN segment to another LAN segment. An effective set of ACL rules should block everything by default and only allow an essential service to flow between network segments after addition of an explicit ACL rule to allow the service to pass. Each rule allowing an additional service (e.g.; file sharing, web services, or FTP) should be considered very carefully to determine whether the service is essential to the effective operation of the ICS network.

Table 1 Suggested Access Control Constraints

Figure 2 shows an ICS network example in which both the SCADA and engineering access are implemented via serial network connections rather than TCP/IP services. Such serial communications architectures dominate older ICS network installations and are still very common in new installations.

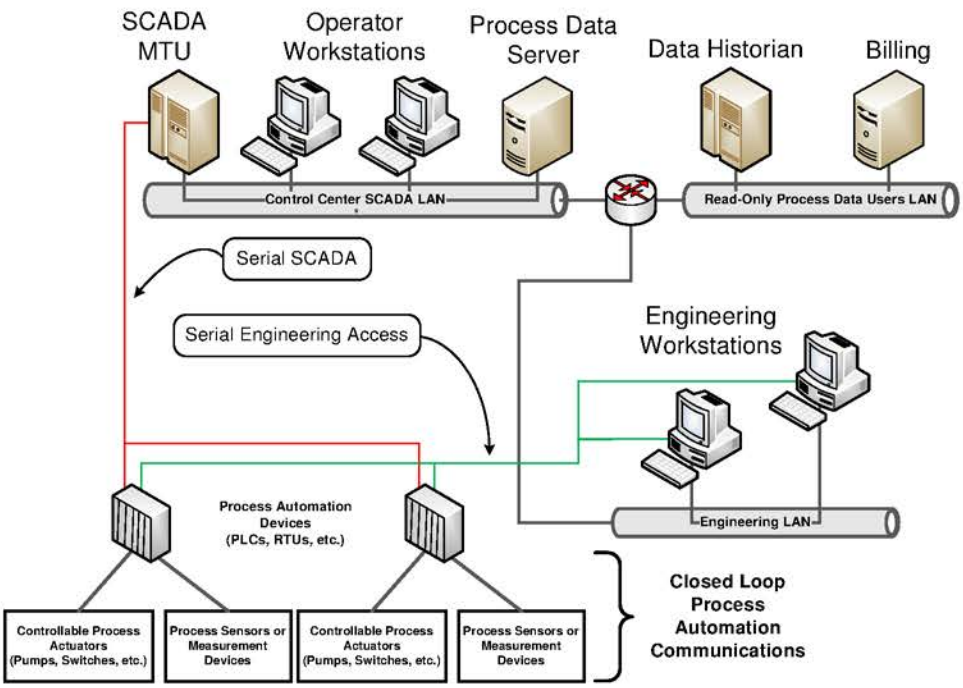


Figure 2 Segmented ICS Network with Serial SCADA and Engineering Access Links

The suggested network architecture shown in Figure 2 also follows the network segregation suggestions discussed previously. The TCP/IP portions of the network are separated into three distinct network segments, isolating the Control Center SCADA assets, Read-Only Process Data Users assets, and the Engineering assets. The process automation devices are not TCP/IP enabled and are, instead, connected to the SCADA and engineering access clients via serial communications links. It is important to note in Figure 2 that the serial engineering and SCADA interfaces are physically separated on each process automation device (e.g.; they are physically separate ports). This ensures that device configuration capabilities are not extended up to the control center SCADA LAN via the SCADA MTU, and that unauthenticated SCADA process control commands are not accessible to workstations on the engineering LAN. Direct serial SCADA connections between the SCADA MTU and process automation devices ensure that these dedicated SCADA services cannot be directly accessed from any other point on the network. Similarly, direct serial engineering access connections ensure that engineering access services are isolated to the connected workstations on the engineering LAN. Finally, the router connecting the three TCP/IP network segments can enforce any necessary traffic filtering suggestions. The first and third ACL rule suggestions in Table 1 are not applicable to the example network in Figure 2; the traffic segregation created by these rules is already present due to the use of direct serial SCADA and engineering access connections. The second and fourth rules are, however, still applicable and should be enforced by the router.

3.2. Authentication

Authentication mechanisms force a user or device to provide proof of identity before receiving access to a networked asset or communications service. Authentication techniques usually require the service user (human or networked device) to prove possession of a protected item or piece of information (e.g.; a

password, cryptographic key, or keycard), or to scan a unique attribute (e.g.; a fingerprint). Ideally, only authorized users can produce the required authentication parameters, so unauthorized access to the asset or service is blocked. In practice, however, the strength of authentication mechanisms common in networked ICS equipment varies a great deal.

3.2.1. User Authentication

The principle of least privilege stipulates that users on the ICS network be granted access only to the equipment and functionality required to conduct their assigned duties. ICS equipment, including PC workstations and embedded logic controllers or automation equipment, often employ password-based login schemes to restrict access to critical functionality. Passwords protect against unauthorized access to critical ICS functionality like programming process control devices or operating pumps, electrical breakers and other sensitive equipment. A strong password-based authentication implementation should support complex password strings, including support for at least eight character passwords consisting of upper and lower case letters, numbers, and special characters (*, -, %, etc.). PC operating systems support strong passwords, but embedded ICS devices often do not. It is important to evaluate the strength of the password schemes implemented in critical ICS equipment to get an accurate assessment of the risks associated with networking each password-protected electronic service. Weak passwords are susceptible to automated password guessing techniques that, if successful, will allow an attacker to gain access to the critical functionality "protected" by the compromised password.

If the equipment supports it, a unique login identity should be created for each user that is to be granted any level of electronic access to the asset. Ensuring that no two users share the same login information enables administrators to accurately log user activity and to hold individuals accountable for their actions. Improved user accountability reduces the risk of malicious electronic attacks by insiders.

The majority of non-PC, embedded ICS devices do not support unique login identities for all users. Such devices are often protected by one or more, configurable password strings that, when entered by a user, grants electronic access to the protected functionality. All authorized users must know the password values and authorized users cannot be differentiated from one another in the activity logs (i.e.; authorized users share login "accounts"). Oftentimes, different usernames and/or passwords are provided to grant access to subsets of the device functionality. For example, a configurable password may authorize read-only access to the device while another password authorizes read/write access. Such authentication schemes are often referred to as role-based authentication because a fixed number of login "accounts" are provided to grant selective access to functionality consistent with a given user's role or job function. When role-based authentication methods are available in an ICS device, authorized users should only be provided the login parameters that grant the least amount of functionality required to conduct their duties. Security policies could be implemented to maintain user level accountability by requiring documentation of which role-based logins were used by which users during what specific time frames.

The table below contains a summary of recommended user authentication practices. These recommendations are supported by almost all PC operating systems, but may not be supported by some software packages and embedded ICS devices. These recommendations should be implemented in all networked assets and software utilities that support them.

Security Recommendation	Cost/Difficulty	Comments
Choose strong password values	Easy	In order to prevent trivial password guessing attacks, password values should be sufficiently long (e.g.; a minimum of eight characters) and contain upper and lower case characters, numbers, and/or special characters. They should not form a word or common acronym.
Assign unique login parameters to every authorized user whenever possible	Easy	Assigning unique username and password values for every authorized user on the system ensures that all users are accountable for actions performed under that account's authorization. Remove users when they no longer require access.
Do not use default password values	Easy	Many embedded ICS devices and software packages ship with password values set to published, default values. These values can often be found in product literature and are easy to guess in password guessing attacks.
Use access privilege settings, and control the distribution of role-based login parameters to grant minimum required access privileges to all users (enforce the principle of least privilege)	Easy	PC operating systems include many security-related settings that can be used to limit the actions that a user can perform while logged in. Many embedded ICS devices and software packages include similar settings that can be used to enforce the principle of least privilege. Consider policy compensations to maintain user accountability in a role-based login environment.

3.2.2. Peer-To-Peer Authentication Services

Devices on ICS networks must communicate with each other to perform critical process control functions. For example, SCADA MTUs periodically retrieve process data from process automation equipment and process data servers transfer and archive information from the SCADA MTU. These automated, device to device communications are not usually protected by password-based authentication mechanisms. Cryptographic authentication services can be used to protect these connections if the devices support such features. For example, newer versions of the Windows operating system support IP Security (IPSec) cryptographic associations natively. PC to PC communications can be protected by setting up an encrypted and authenticated TCP/IP communications "tunnel" between them. In addition, network domain settings can be employed to authenticate PC to PC connections.

Support for cryptographic authentication is still limited in most embedded ICS devices. If a device to device communications link is considered especially insecure or risky, inline cryptographic devices can be employed to augment the security of the link. Various products exist to secure serial or TCP/IP links. For example, virtual private network (VPN) products based on the IPSec or Transport Layer Security (SSL/TLS) protocols are common, fairly inexpensive, and extremely effective at securing TCP/IP

network links. Products are also available that are designed to be compatible with serial process control protocols like those carried on multi-drop or point to point serial SCADA networks.

3.3. Cryptographic Security and Data Confidentiality

Passwords, sensitive files (e.g.; network architecture diagrams), or other private data may be susceptible to interception during transmission over network links. Allowing sensitive data to traverses network segments (serial or TCP/IP) accessible to users not authorized to view the data increases the potential for insider attacks. Connections to multi-drop serial networks or continued use of TCP/IP hubs[3] may put a motivated attacker (insider or outsider) in a position to capture sensitive data directly from the network. Furthermore, an attacker may tap network media or redirect TCP/IP frames on a switched network segment (usually by manipulating the Address Resolution Protocol caches on network switches) to allow interception of data that they would otherwise not have access to.

Many remote access services being used in ICS devices (e.g.; for networked engineering access) transmit user authentication parameters in a format that can be trivially extracted from captured network data. Captured passwords can be used by an attacker to gain unauthorized access to critical ICS devices. Many serial and TCP/IP Telnet terminal applications like HyperTerminal transfer all data, including sensitive login passwords, in simple ASCII character format that can be read directly from intercepted data. The FTP and TFTP file transfer protocols transmit all requested files unencrypted. FTP even transmits login passwords unencrypted (TFTP does not support password protection). Finally, most third party engineering access software suites, supplied by ICS equipment vendors for use on their products, do not include features to obfuscate transmitted device login passwords.

The following recommendations can be used to reduce the risk of data interception.

Security Recommendation	Cost / Difficulty	Comments
Replace TCP/IP hubs with switches	Easy	Switches perform better and are more secure than hubs. Removing hubs should be considered a routine network upgrade.
Prevent physical access to network media and unused network access points	Easy to Moderate	Many physical security upgrades can be installed without disrupting the network. Upgrades that involve network downtime and equipment relocation can be difficult.
Use static Address Resolution Protocol (ARP) tables in network switches	Easy to Moderate	Switches that support static ARP tables are inexpensive and widely-available, but manually managing ARP table updates can be inconvenient on anything but highly static networks.
Use inline encryption devices or peer-to-peer encryption services to protect data in transit	Moderate to Difficult	PC-based ICS equipment may support cryptographic services, but embedded ICS devices usually cannot be augmented with additional features without extensive firmware modifications by the vendor. Inline cryptographic modules can be difficult to install and maintain.

[3] Older TCP/IP hubs replicate all received data to all physical ports on the device, thus, all received frames are visible to all devices connected to the hub. Switches and routers, on the other hand, forward received frames only to the physical port to which the intended recipient is connected.

The first three recommendations reduce a potential attacker's ability to gain access to data transmitted over the network and the last entry uses cryptographic services and/or devices to encrypt the data to prevent intercepted data from being interpreted by unauthorized individuals. Some of these recommendations may need clarification.

Many TCP/IP network switches can be operated in a mode in which ARP table entries are entered manually instead of being dynamically managed by the switch itself. ARP tables simply map a host's Media Access Control (MAC) address to the IP address assigned to its TCP/IP network interface. Non-static ARP tables can be manipulated by attackers to reroute TCP/IP traffic. Static ARP tables help prevent these types of data routing attacks. Static ARP tables are fairly easy to maintain on network segments on which new hosts are seldom added and IP addresses seldom change, but can be hard to manage in more dynamic environments. Fortunately ICS networks are typically fairly static.

We outlined some remote access and file transfer services that do not encrypt transmitted data. Sometimes these services can be replaced with more secure services that protect data in transit. For example, the Secure Shell (SSH) service can be used to implement a cryptographically secure terminal and file transfer capability. PC-based workstations and ICS equipment can often be augmented with additional secure software and services. Unfortunately, in embedded ICS devices we are largely stuck with the remote access services and software that the vendor has implemented in the product. Older, and many current embedded ICS devices do not implement secure remote access services that encrypt data prior to transmission and cannot be directly upgraded to support such services due to lack of vendor support for such upgrade paths. It is for this reason that we recommend isolating engineering workstations that communicate directly with embedded ICS equipment as shown in **Figure 1** and 2. If necessary, inline cryptographic devices can be added to secure connections to embedded ICS devices that are deemed exceedingly risky. We discussed some of the available inline cryptographic devices in the authentication section above. Most of the cryptographic devices that provide strong authentication services also provide encryption capabilities to provide effective data confidentiality.

3.4. Device Hardening

PC workstations and servers, as well as embedded ICS devices often support configurable options that can improve the security posture of the asset. This is especially true for PC operating systems. Security policy settings exist in most PC operating systems that allow administrators to customize security-related features. By default, newly-installed PC operating systems are configured to strike a balance between ease of use and security. These default configurations can, and should, be customized to ensure that commissioned PCs are adequately secured. Extensive PC hardening guidelines are available from the NSA.gov website. These publications are excellent guidelines but should be used cautiously in an ICS environment. Some of the recommended security settings may reduce PC availability below acceptable levels. For example, some publications recommend that user passwords automatically expire every few months. When the PC is used regularly, the user will receive adequate notification that the password value must be changed. If the user misses or ignores the password change prompts and the password expires, they will not be able to log in until a system administrator resets the password. The risk of being locked out of critical ICS workstations or servers during an emergency may be unacceptable.

Embedded ICS devices often include some of the security-related configuration options common to PC operating systems. Below is a table of recommendations that summarize the use of some of the features common to PCs and embedded ICS devices.

Security Recommendation	Cost/Difficulty	Comments
Turn off all unused services and disable unused communications ports	Easy	Electronic attacks often exploit flaws in vulnerable communications services. Disabling unused services or communications ports eliminates some of these potential attack vectors.
Enable strong password enforcement	Easy	Strong passwords should always be used. Most PC operating systems support security settings that force the user to enter a sufficiently strong password (e.g.; minimum number of characters and forced use of a diverse character set). Support for strong password enforcement is still minimal in most embedded ICS devices.
Enable failed login timeouts	Easy	Failed login timeouts will lock out user access if a preset number of consecutive login failures occur (e.g.; wrong password values are entered by the user). Enabling failed login timeouts on PCs or embedded ICS equipment greatly reduces the rate of automated or manual password guessing attacks. However, be cautious when enabling this feature in critical ICS equipment to avoid excessive loss of system availability. A short 30 second timeout after 5 consecutive failed logins is still sufficient to slow password guessing attacks while ensuring adequate system availability.
Follow PC operating system hardening guidelines (e.g.; those published by NSA etc.) whenever feasible	Easy	Be cautious when applying PC recommended hardening practices to avoid excessive loss of system availability.

3.5. Security-Related Firmware and Software Updates

Many electronic attack techniques exploit firmware and software flaws to negatively affect targeted systems. Some of the most damaging viruses and worms seen to date utilize flaws in communications services to self replicate and spread. Electronic attacks can come from both accidental introduction of malicious viruses and worms onto the ICS network, or opportunistic insiders or outsiders conducting active attacks on ICS network assets. The effects of such attacks often include forced system crashes, unauthorized access to data, or even complete target takeover. Buffer overflow flaws launched remotely against a vulnerable communications service, for example, may allow an attacker to execute malicious code on the target system. The executed code may include instructions that allow the attacker to remotely log onto the target system with the highest authorized privilege (e.g.; administrator or root).

The network segmentation recommendations discussed earlier are meant to reduce the number of potentially vulnerable services available at various points on the network. If a service flowing from one network segment to another contains flaws or weak authentication mechanisms that make them vulnerable to electronic attack, an attacker may be able to take over the system hosting the vulnerable service to propagate the attack across network segments. Referring to **Figure 1** for example, if ACLs were not correctly implemented, an attacker on the 'Read-Only Process Data Users LAN' segment may discover a flaw in the database service running on the process data server on the 'Control Center SCADA LAN' that allows the attacker to execute malicious code. The executed code could install a backdoor terminal service on the process data server that communicates on a TCP/IP port that is not blocked by the router. The attacker is now in a position to launch new attacks from the process data server against any target on the 'Control Center SCADA LAN'. Weak authentication mechanisms in the process control protocols transferring process data and commands between the SCADA MTU and the operator workstations could allow the attacker to issue process control commands from the compromised process data server.

Vendors often release patches for PC operating systems and third-party software to eliminate discovered security flaws. Security-related firmware patch releases for embedded ICS devices are less common simply because they are not subjected to vulnerability-revealing attacks as often as PC operating systems and software packages are. This is a direct consequence of the fact that PCs represent the vast majority of targeted systems for Internet-born electronic attacks.

It is extremely important to apply high priority, security-related software and firmware patches as soon as possible to remove known vulnerabilities before they get exploited in an active electronic attack. Most patch releases include extensive explanations of which services or features are being fixed and the consequences of attacks exploiting the discovered flaws. Patches that fix high risk flaws (e.g.; ones that allow remote execution of malicious code) in services that are actively used and potentially exploitable on your ICS network should be given high priority. Patches that fix low risk flaws, provide simple user interface or feature improvements, or fix flaws in services that are not used on your ICS network can be considered low priority.

Applying patches may have unforeseen consequences on a commissioned network. Services may break, PC operating systems may crash or become unbootable, and software packages may no longer work after the patch is installed. The risk of these unforeseen consequences may not be acceptable in an ICS network environment, especially for critical servers, workstations, and embedded ICS devices. Because of this, patches should never be applied to critical ICS assets without thorough testing. Patches should be tested offline on a test network, to ensure compatibility and system stability prior to applying the patch to active ICS network assets. Vendor supplied, ICS-specific software packages like human-machine interface (HMI) software, process control bus interface drivers, engineering access suites, or process data management suites can be particularly susceptible to unforeseen consequences after installing operating system patches. Patches to update ICS vendor-supplied software packages may significantly lag publication of the operating system patch that fixes them. If offline tests indicate that critical software packages are incompatible with a high-priority patch, you will have to contact the software vendor to see if a fix is available, or insist that a fix be made available. Most patches require a hard reboot of the device after installation, so tested patches must only be applied to devices that can safely be rebooted

while the process control system is online, or at times when the process control system can be brought offline or put in an otherwise safe state.

It is extremely important to realize that PC operating systems evolve with time and become less vulnerable due to the release of new patches and operating system versions. Operating system vendors will stop supporting old versions at some point in time and security-improving patches will no longer be released for that version. Some older, unsupported operating system versions like Windows 98 or Windows NT have well-known, severe security flaws that will never be fully removed. All PCs in use on ICS networks should be upgraded to operating system versions that are still being supported and patched by the vendor. If testing shows that software packages running on the PC conflict with the new operating system version, then they should also be upgraded to remove compatibility conflicts. Any PCs that cannot be upgraded to supported operating system versions due to irresolvable compatibility conflicts should be isolated on their own dedicated network segment. All but the most essential services to and from the segment should be blocked by a router or firewall, especially all operating system services that are known to have irreparable vulnerabilities.

3.6. Security Event Monitoring

PC operating systems include extensive event logging features that allow administrators to monitor user activity on workstations and servers on the ICS network. Log entries record events like successful and failed user logins, password changes, program and file access, and much more. Embedded ICS devices often include similar features that can be used to track device configuration changes, execution of process control commands, and other user activities. When these features are available, they should be enabled and security-related event logs should be regularly monitored to increase user accountability, reduce insider threats, and potentially detect the precursors of an electronic attack. Where possible centralize all logging for easier analysis and use a common time source for all devices so logs will be synchronized.

4. Preserving ICS Performance and Functionality

Reliable process control and monitoring are essential for the safe and stable operation of industrial control systems. Some security solutions have the potential to disrupt the availability of critical control system functionality. Such disruptions can hinder situational awareness, degrade the ability to apply corrective user inputs, or otherwise compromise the integrity of the process. Application of defensive security solutions must not degrade critical ICS functionality beyond acceptable levels.

4.1. Bandwidth Limitations

Some ICS network links are implemented using older, low-bandwidth communications technologies. Low baud rate serial modems (e.g.; 1200 baud Bell 202T) or network interface converters (e.g.; EIA-232 to EIA-485 converters) can cause network bottlenecks that may not accommodate certain security technologies. A heavily-taxed, bandwidth-limited SCADA connection, for example, may not function reliably after installing inline cryptographic modules. Cryptographic protocols that are not designed for such environments may introduce excessive data overhead due to insertion of cryptographic headers, authentication hash bytes (often upwards of 8 to 16 bytes per transmitted message), or protocol messages

associated with maintaining the link. The amount of overhead added by cryptographic modules varies with the brand and model. Some models are available that minimize data overhead and are therefore less likely to negatively affect the link on which they are installed. Bandwidth considerations must be addressed before installing security-related products on critical ICS network links.

4.2. Latency Requirements

Some ICS network links are highly susceptible to message delivery delays. Closed loop process control logic can often run with as little as tens of milliseconds between execution cycles. Communications between the process sensors/actuators and the process automation devices running the control logic cannot accommodate excessive transmission latency. Because of this, it is usually unwise to insert additional security devices on closed loop process control network links. Keeping the process automation devices in close proximity to the process sensors/actuators to which they are connected and ensuring that the sensors/actuators are sufficiently isolated from remote network segments (see **Figure 1**and **Figure 2**) should be sufficient to meet the latency requirements of most ICS networks.

Some process control and SCADA protocols cannot tolerate disruption of the timing between transmitted bytes (inter-character spacing). Modbus protocol frames, for example do not include header fields that a receiving host can use to reliably identify individual frame boundaries within the byte stream. Instead, most Modbus interfaces simply rely on inter-character delays to identify frame boundaries: bytes within a given frame must not be separated by too much time, while consecutive frame transmissions must be separated by a sufficient amount of time. Inline security products that alter these strict timing rules will severely disrupt the reliable reception of message frames. There are cryptographic products available that are specifically designed for such environments, but most are not. Care must be taken to identify network links that have strict timing requirements and to ensure that employed security products are compatible with these requirements.

4.3. Availability Requirements

Critical ICS equipment must remain functional and accessible to other critical equipment and to authorized users. The addition of security enhancing devices and features, including cryptographic modules or services, routers, firewalls, and strong authentication mechanisms, will add additional complexity and points of failure on the ICS network. For example a mistake made while reconfiguring a security device (e.g.; a syntax error in a new router ACL entry, or a mismatch in cryptographic keys) may stop a critical communications service until the mistake is discovered and repaired. As mentioned previously, access lockouts due to password expiration or failed login timeouts can prevent authorized user access at critical times: it is quite conceivable that an authorized user might enter the wrong password several times and accidentally lock themselves out in a time of stress or emergency – exactly when user access is needed the most.

In most cases the potential for availability reduction is worth the increased security that these products and techniques provide. It is, however, important to consider the availability impacts when choosing security enhancements. It is almost always possible to strike a balance between adequate security and maximized availability.

5. Conclusions

Most ICSs in operation today were not designed with today's threat vectors in mind. Closed ICS networks, by definition, are not susceptible to "outsider" network threat vectors. However, there are many other threat vectors which are of concern for asset owners. This paper has offered a variety of suggestions for asset owners to improve the security posture of their system while preserving the functional capability and availability necessary to fulfill the system Concept of Operation (CONOP).