May/June 2014

InTech®

A PUBLICATION OF THE INTERNATIONAL SOCIETY OF AUTOMATION

Integrating DCS I/O
Embedded vision
Multigenerational systems
Mobile user interfaces
Flow spotlight

Top ten differences between ICS and IT cybersecurity

Understanding leads to cooperation and collaboration between historically disconnected camps

Setting the Standard for Automation™

www.isa.org/intech

Top ten differences between ICS and IT cybersecurity

Understanding the different needs of ICS and IT system security leads to cooperation and collaboration between historically disconnected camps

By Lee Neitzel and Bob Huba

In many, if not most plants with industrial control systems (ICSs), ICS engineers and their internal information technology (IT) counterparts have very different perspectives on cybersecurity. Not surprisingly, these different perspectives often lead to conflicts when connecting an ICS to the plant's IT system.

In the past, because ICSs used proprietary hardware and software, this interconnection focused primarily on just being able to communicate. The introduction of Ethernet and Microsoft Windows into ICSs in the mid-1990s, followed by the development of OPC interfaces, greatly simplified this problem, but at the cost of exposing the ICS to security threats previously known only to IT systems.

Further, with the rapid increase of attacks on industrial systems in the past few years, chief information officers are often held responsible for cybersecurity for the entire plant, including their ICSs. Unfortunately, not all IT security solutions are suitable for ICSs because of fundamental differences between ICS and IT systems. In addition, plants often have multiple production processes and ICSs, and some are naturally more critical than others. As a result, it is not uncommon for security to be handled differently among the various ICSs in a plant.

This article discusses how ICSs differ from IT systems as they relate to cybersecurity. It is important that IT and ICS professionals jointly understand the following top ten differences and develop workable security solutions that benefit the whole organization.

Difference #1: Security objectives

One of the biggest differences between ICS and plant IT security is the main security objective of each. Plant IT systems are business systems whose primary cybersecurity objective is to protect data (confidentiality). In contrast, the main cybersecurity objective of an ICS is to maintain the integrity of its production process and the availability of its components. Protection of information is still important, but loss of production translates into an immediate loss of income. Examples of threats to production integrity include those that degrade production, cause loss of view/control, damage production equipment, or result in possible safety issues.

One of the consequences of ICSs focusing on the production process is that ICS security is implemented using a comprehensive set of defense-in-depth layers to isolate the ICS and the physical process from the plant IT system. This isolation is the topic of difference #2.

Difference #2: Network segmentation

The first difference encountered when connecting ICS and IT systems is how they are segmented and protected. IT systems are usually composed of interconnected subnets (short for "subnetworks") with some level of Internet connectivity. As a result, access controls and protection from the Internet is a primary focus of IT network security. It is not uncommon to see sophisticated firewalls, proxy servers, intrusion detection/prevention devices, and other protective mechanisms at the boundary with the Internet.

Inside this boundary, the remainder of the IT network is segmented into subnets that are generally aligned with organizational and geographical boundaries. Because access between these subnets is usually required, security between them is typically limited. However, all traffic from them must pass through the Internet security boundary to access the Internet. ICS networks, on the other hand, can be viewed as industrial intranets with two overriding security requirements. First, no access to the Internet or to email should be allowed from ICS networks. Second, ICS networks should be rigorously defended from other plant networks, especially those with Internet access.

To meet these requirements, ICSs usually employ network security devices (e.g., firewalls) for isolation from the plant IT system. Only workstations and servers within the ICS that act as gateways should allow ICS access through these ICS perimeter security devices. This prevents other devices on the ICS control network from being directly accessible from the plant network. These gateways should have an additional network card that allows them to connect the ICS control network. In general, only devices authorized to access the ICS from the plant network should be aware of these ICS network security devices and therefore be able to send messages through them to ICS gateways.

ICSs should be further insulated from the plant IT system by a demilitarized zone (DMZ) that sits between the plant network and the ICS. The DMZ is an intranet that should be hidden from the plant network by an undiscoverable network security device. All external access to the ICS should first pass through this device and then be terminated in DMZ servers. DMZ servers provide clients on the plant network with ICS data and events that these servers independently obtain through separate and isolated communications with the ICS. The network security device that connects the DMZ to the ICS should be configured to allow only these isolated communications to ensure that all ICS access goes through the DMZ servers.

As a further precaution, the DMZ should use private subnet addresses that are independent of subnet addresses used in the plant network to prevent plant network messages from being erroneously routed to the DMZ. Similarly, the ICS should use private subnet addresses that are independent of DMZ addresses.

ICS networks often have remote input/output (I/O) systems, whereas IT networks do not. In these systems, I/O devices are installed in remote geographical locations and are often connected to the ICS via modems over public networks, virtual public networks (VPNs), and satellite links. Care must be taken, because these connections can give rise to security issues.

FAST FORWARD

- Differences in ICS and IT security objectives cause competing and often conflicting security solutions.
- Differences in ICS and IT system characteristics lead to different defense-in-depth strategies.
- Differences in ICS and IT operational characteristics cause differences in how security mechanisms are implemented and used.

Difference #3: Network topology

Closely related to network segmentation differences are network topology differences. Many IT systems are large when compared to a typical ICS and contain data centers, intranets, and Wi-Fi networks. ICSs, on the other hand, are often small and have only a configuration database and data/event historians.

It is not uncommon for an IT system to have hundreds if not thousands of nodes whose numbers change daily as employees come and go, as applications evolve, and as mobile devices are connected and disconnected. In contrast, most ICSs are an order of magnitude smaller, and generally have statically defined configurations.

IT network configurations, including VPNs, and network security devices have to keep up with these changes. As a result, IT systems extensively use many automated tools, such as dynamic host configuration protocol (DHCP), to manage their network topologies. These and other tools are cost effective only in large-scale systems and are considered expensive and complex by ICS standards.

ICSs typically remain relatively static for years. A rigorous change management process is normally mandatory to ensure all changes are approved and tested. In addition, the use of DHCP and Wi-Fi segments are discouraged in the ICS for security reasons. In addition, ICS networks that connect ICS workstations with controller-level devices are normally redundant to prevent a network failure from affecting the operation of the control system. This network redundancy is typically proprietary to the ICS vendor with custom addressing models and switchover logic. As a result, the tools and techniques

Those responsible for cybersecurity within an organization must understand the differences between ICS and IT systems in order to work together effectively.

IT uses to maintain its dynamic network topologies are often not suitable or applicable to statically defined ICS networks.

Difference #4: Functional partitioning

ICS and IT systems are functionally partitioned in different ways. The most common approach taken by IT systems is to divide the system into various administrative partitions to better restrict user access to information assets. The IT department typically implements the partitions using Windows Domains and operating system objects, such as files.

Domains and organizational units typically represent business units/geographical entities within an organization, to which users and computers are assigned. Groups are used to control access to these computers and their objects (files, folders, executables, etc.) through the definition of access control lists (ACLs).

Each object contains an ACL that identifies who has been granted/denied access to the object. To simplify the process of pairing users with objects, groups are defined and assigned to objects, and then users are assigned to groups. As a result, only users/roles who are trusted to access an object are granted permission to do so. The careful definition of groups/roles can thereby be used to partition an IT system into trust levels.

ICS partitioning is much different. The ICS is partitioned into three levels (0, 1, and 2), as defined by the ISA95/Purdue reference model. Level 0 represents the physical process; Level 1 is control and monitoring; and Level 2 is supervisory control. Because of the nature of the devices used in these ICS levels, it is necessary to map trust levels to the device. In this case, trust means how much a device is trusted to behave as expected.

At Level 1, field devices perform I/O operations on the physical process (Level 0). Because they operate on the physical process, field devices have the highest level of trust. Trust generally is ascertained through design reviews, functional testing, and experience. Devices whose behavior is questionable should not be trusted and should not be used in Level 1.

Field devices use proprietary designs and firmware. Many can communicate digitally using standard, industrial protocols such as HART, Foundation Fieldbus, Profibus, DeviceNet, and Modbus. With the exception of wireless, field device protocols rarely include security features. Therefore, access to field devices must be protected by external means. Unfortunately, network security devices, such as firewalls, that are commonly used in IT systems are not applicable. These industrial protocols are not based on Ethernet or TCP/IP. Instead, physical and procedural security often restricts access to field devices and their communication links.

In addition, device firmware needs protection, including protection of upgrade files and the processes used to install them (e.g., flash upgrades and over-the-wire upgrades). Currently, the firmware upgrade process often has limited security features.

At Level 2 are distributed control system controllers, programmable logic controllers, remote terminal units (RTUs), remote I/O devices, and other similar devices. Because they read and write field device parameters, controller-level devices require the second highest level of trust, generally attained through testing and experience.

Controller-level devices, other than some RTUs and other remote devices, usually have limited security-related features and rely on the Level 2 control network for protection. ICS vendors often use industrial grade, proprietary firewalls and Ethernet switches in the control network to separate it into two layers, the workstation layer and the control layer.

These network devices have three primary security objectives: to lock down the network to prevent unauthorized devices from connecting to it, to protect controller-level devices from unauthorized contact, and to prevent them from being saturated with network traffic by rate-controlling the network traffic flowing to them.

IT typically does not have the policies, procedures, tools, and expertise in place to manage the ICS vendor-specific Level 2 network and controller-level devices and the Level 1 I/O devices.

Also at Level 2, and sitting above controller-level devices, are the workstations/servers—configuration/engineering, maintenance, operator, historian stations—all having direct connectivity to the controllers, and all using components and operating systems familiar to IT, such as PCs, Windows, and Ethernet. Level 2 workstations and servers have the third highest level of trustworthiness in the ICS. They provide the buffer between the outside world (Level 3 and beyond) and the process, so outside direct access to controller-level devices should not be allowed. Access to controller-level devices

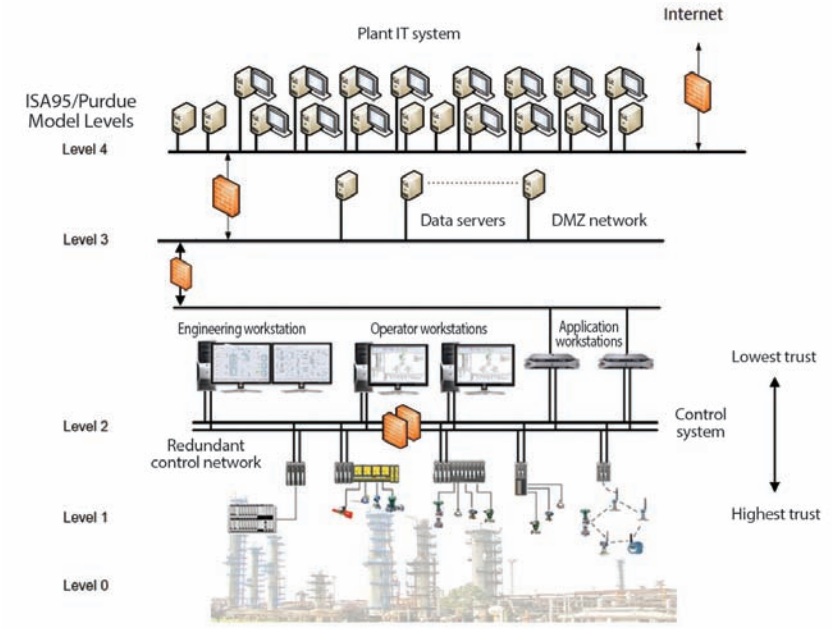

Compared to a typical IT system, most ICSs contain relatively few workstations and other computing components, a crucial difference that greatly affects the feasibility of implementing certain cybersecurity measures.

should be limited to Level 2 workstations and servers approved by the ICS vendor.

The trust levels of Level 2 workstations and servers are lower than controller-level and field devices for three reasons:

- They run commercial operating systems and software (e.g., SQL database software) with vulnerabilities that are continuously being discovered and exploited.
- They have a better chance of being infected or compromised, because they can be accessed by Level 3.
- They have users who may not always follow policies and procedures—some may plug in nonverified USB sticks, plug in their smartphones to charge, or bring in their own software that has not been tested to operate correctly with the ICS.

The trust levels associated with field devices, controller-level devices, and workstations are inherent to most control systems. Understanding them and maintaining separation/isolation between them is a responsibility that is normally not present in IT systems.

Difference #5: Physical components

Closely related to functional partitioning and trust levels are the physical components used to implement ICS and IT systems. IT systems are primarily composed of off-the-shelf networks, workstations, and servers that IT can access and administer. As a result, IT departments are able to define security policies for these components and enforce them with off-the-shelf security-related applications and devices, such as firewalls, antivirus systems, and patch management systems.

In contrast, ICSs are not IT systems doing control, as it may sometimes appear, but instead are tightly integrated proprietary systems. With the exception of workstations and servers, ICSs are composed of components that are generally custom built and foreign to IT. This often includes network devices built for industrial use, including Ethernet switches and firewalls. And, although ICS workstations and servers are typically based on Windows, they are usually hardened by the ICS vendor to the point that their software, other than the operating system, is custom built, and their security policies are set to industry standards that may conflict with the policies used within the IT system.

Consequently, IT security cannot just be mapped onto the ICS. Instead, the components used in the ICS may, and often do, require security-related ICS vendor-specific tools unknown to IT systems, such as custom event logs, port lockdown mechanisms, and features for disabling USB ports.

Difference #6: User accounts

IT systems generally support two levels of users: users known to the operating system (e.g., Windows users) and users of specific applications (e.g., order-entry systems). Operating system user accounts are used to authenticate the user during login and to identify which operating system resources the user can access. IT system administrators often administer operating system user accounts with Windows Domains/Active Directory. When multiple domains are present, IT administration establishes trusts between specific domains to let users access resources across domain boundaries.

IT systems also often contain applications, such as database applications, that have their own user accounts that can be independent of operating system accounts. For these applications, the user must go through a separate login screen before being allowed to access the data.

ICSs also use operating system user accounts and domains. However, allowing IT systems users to access the ICS by establishing trusts from IT system domains to the ICS domain is generally not recommended, since it reduces isolation of the ICS.

ICSs also have their own application-specific users. Unlike IT applications, however, the ICS is really a complete distributed system composed of configuration, operation, and maintenance applications, databases, and event journals. ICSs almost always use role-based access controls for granting/denying access to control data and devices. Operators, process engineers, and maintenance engineers are examples of these roles.

To manage access to these elements of the ICS, ICSs typically have an ICS-specific user management application. Although in principle this is similar to IT application security, the complexity, scope, and technical expertise required to administer ICS users is closely related to the nature of the process being controlled, which is generally not familiar to IT system administrators.

Finally, authorizing access from the plant network to the ICS becomes more difficult because of these differences. Do all external users become users of the ICS and its domain, or do DMZ server applications provide access to authorized IT system users but connect to the ICS using ICS credentials? Also, how is traceability maintained for auditable ICS transactions? Answering these questions normally requires collaboration between the ICS and IT systems administrators.

Difference #7: SIS

Plant safety is a critical part of plant operation, and ICSs, therefore, often include integrated, yet distinct, safety instrumented systems (SISs). The SIS is responsible for maintaining the safe operation of the process by placing the process into a safe state when process conditions that threaten safety are detected. IT systems have no systems analogous to the SIS.

SIS networks are usually proprietary and must be securely segmented and isolated from ICS networks. In addition, the SIS decision-making component, commonly called the logic solver, is also a custom, proprietary component, separate even from other components used in the ICS. Also, SIS-specific standards that include security are currently under development in ISA84. As a result, commonly used IT tools and network devices are not applicable to SIS network security.

Managing the security of an ICS includes an often manual effort to ensure that the SIS is protected from the ICS and from external interference, and that its integrity has not been compromised. These are capabilities not normally within the scope of IT systems professionals.

Difference #8: Untested software

IT systems are typically open systems, which allow them to run off-the-shelf software and to evolve over time. Evolution includes adding new software; updating workstation, server, and network device hardware and software; replacing components as needed; and even adding new components to the system. Keeping systems current is one of the approaches taken in IT systems to maintain security.

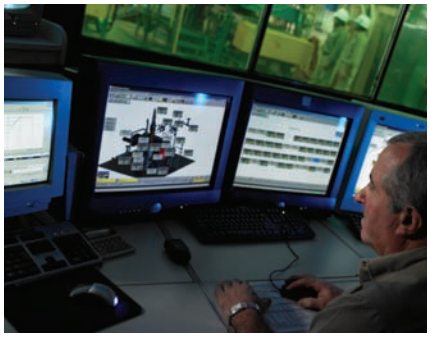

Unlike their IT counterparts, ICS users need additional role-based access controls so that each person can access only the areas of the ICS needed to do a particular job.

ICSs, however, are typically closed and implemented to a specific hardware configuration and operating system version (e.g., service pack), and may not run properly if either is changed. As a result, all updates, including patches and virus definition files, have to be thoroughly tested with the ICS before being approved for installation.

Likewise, all new software added to the ICS that is not supplied or supported by the vendor should be thoroughly tested for compatibility with the ICS. In some cases, as with those regulated by the Food and Drug Administration, the ICS and IT systems associated with the regulated product must be validated, and once validated, cannot be updated with new software without being revalidated. But for typical IT systems, this rigor is not common. Running software that has not been tested with the specific ICS is a serious concern, because of its potential to cause conflicts or failures within the ICS or introduce vulnerabilities of its own. Therefore, all software to be run in an ICS should be tested and approved using a formal operations change management process.

The most common way to protect against the introduction of unapproved software is to restrict installation privileges and to use access control lists for program directories. However, these mechanisms do not protect against executables that can be copied to the directory and run without being installed. Mechanisms to prevent this type of software from being loaded onto a workstation include disabling USB ports and CD/DVD drives and tight control or elimination of shared drives. Although these are commonly employed techniques in ICS workstations and servers, they are seldom used in IT systems.

Mechanisms to prevent unapproved software from being run are not as commonplace. While antivirus software can detect infected software, it cannot detect untested or unapproved software. For this, whitelisting is gaining acceptance in IT systems. Whitelisting complements antivirus programs by allowing only approved and authentic (uninfected) executables to run. However, because of the checks necessary to validate an executable each time it is run, performance is affected.

Software that has been approved to execute in an IT system often has not been rigorously tested for compatibility with the IT system. All software that is allowed to run on an ICS must be tested to ensure it will not interfere with the ICS.

Difference #9: Patching

IT systems normally have patch management software that automatically installs security updates very quickly after their release. On the other hand, it is not uncommon for patches to be deferred or postponed indefinitely in ICSs. ICS patching requires testing, approval, scheduling, and validation to ensure safe and repeatable control. Scheduling is required because of the potential disruption to operations, such as reboots. Reboots can cause a temporary loss of view/control, and worse, they can fail, often requiring technical intervention to return a failed component to service. As a result of the effort required and because of the associated risks, patching is often not performed on an operational ICS, or at least not on the same schedule as IT system patching.

In addition, because the lifespan of ICSs is so long, patches for many older systems are no longer available. For example, there are many ICSs still in operation that run Windows NT and Windows XP.

The challenge for ICSs, which is not shared by IT systems, is to keep unpatched systems secure. Typically this is done through compensating security mechanisms in an ICS's defense-in-depth strategy.

Difference #10: Security inconveniences

As most of us probably agree, cybersecurity measures add a degree of inconvenience to our jobs. Who has not had to wait while operating system patches are being installed? Or who has not had to call the service desk to report that he or she is locked out and needs to have a password reset? But as cumbersome as they can be, we have all learned to live with these inconveniences.

However, in an ICS environment, such inconveniences may not be tolerable, especially those that decrease performance. Imagine not receiving a critical system alarm in time to respond to it, or having to handle it while the workstation decides to reboot itself. Also, having to use a long and complex password during a process upset may not be acceptable. While many of these inconveniences are not specific to ICSs, they can be intolerable to them.

As a result, security measures that are acceptable in IT systems may not be acceptable in an ICS. If indiscriminately employed in an ICS, IT security measures may pose one of the biggest threats to ICS security. Because they are so painful or disruptive, they often result in the security mechanisms being bypassed, disabled, postponed, or otherwise ignored. Not only will this expose the ICS to vulnerabilities, but it will also negatively affect attitudes of ICS users toward future attempts to secure the ICS.

We have examined how ICSs differ from IT systems with respect to cybersecurity. Unfortunately, failure to understand these differences often leads to conflicts between IT and ICS administrators, which leads to a less-than-optimal security solution for the plant. These discussion points should help promote communications and resolve conflicts. ■

ABOUT THE AUTHORS

Lee Neitzel (Lee.Neitzel@Emerson.com), senior engineer at Emerson Process Management, has been involved in security and network standards for more than 25 years. He is currently the IEC project leader for integrating the WIB "Process Control Domain – Security Requirements for Vendors" specification into the ISA-99/IEC 62443 security standards. **Bob Huba** (Bob.Huba@Emerson.com), system security architect, has been with Emerson Process Management for 36 years. He is active in the development of the ISA-99/IEC 62443 standards.

View the online version at www.isa.org/intech/20140601.