

White Paper

Control System Cyber Vulnerabilities and Potential Mitigation of Risk for Utilities

By Joe Weiss PE,
CISM Applied Control Solutions, LLC



With contributions from Manoj Apte PhD,
Juniper Networks



Table of Contents

- Executive Summary 3
- Computer Networking in Today’s Industrial Facilities 3
- The Security Challenges of Control Systems 4
- Threats to the Control Network 5
- The Impact of Security Incidents on Industrial Control Systems 7
- Meeting New Regulations for Network Security 8
- Typical Security Technologies Applicable to Control System Networks 10
 - Firewalls 10
 - Intrusion Detection and Protection 10
 - Authentication/Authorization Systems 11
 - Network Access Control 11
 - Secure Remote Access 12
 - Encryption of Critical Data 12
 - Monitoring for Administration and the Audit Trail 12
 - Configuration Management 12
 - The Need for Control System Security 12
- About the Author 13

Executive Summary

Today's global industrial infrastructure includes thousands of electric utilities, water/wastewater management companies, oil and gas suppliers, chemical manufacturers and other facilities critical to daily functioning. At the same time, the control and monitoring system networks that operate these infrastructures are among the most vulnerable targets for hackers, extortionists and disgruntled or ineffective employees. As organizations progressively rely more on off-the-shelf operating systems and Internet-based remote access to carry out production tasks, traditional control networks are open to the worms, viruses and application-level attacks that proliferate online. One means of securing these vulnerable systems, without replacing the entire industrial infrastructure, is to use network-based security.

Computer Networking in Today's Industrial Facilities

Today's industrial facilities are increasingly moving toward computer networks that integrate previously isolated control systems—for electric power, oil and gas, water and similar services—with corporate IT and other systems. By connecting all of the organization's computer systems based on open networking standards, they are able to operate more efficiently and effectively. This allows plant managers to integrate sensor and enterprise data to increase visibility throughout the facility, helping to keep systems available 24 hours a day, assisting in rapid resolution of problems, and reducing operational and support costs. For example, based on open network access:

- Substations and pump-houses located in remote regions can be centrally controlled, monitored and maintained at a lower cost.
- The data gathering and audit-report generation required for regulatory compliance is accomplished more easily and effectively.
- Industrial end users, including power, oil and natural gas producers (ONGs), have instantaneous access to current plant data, from the specific gravity of oil flowing in a pipeline to the amount of power being generated by a plant.
- Energy brokers can trade commodities based on real-time production numbers, saving billions of dollars for utilities and ONGs.
- Electric utilities can provide required data to Independent Service Operators (ISOs).
- Other industrial manufacturers, such as those in chemicals and auto manufacturing, can optimize production with just-in-time deliveries.
- Online video surveillance and lower-cost voice over IP (VoIP) telephone technologies can help secure and maintain facilities and plants.

Among the control systems connected to these new integrated networks are Supervisory Control and Data Acquisition (SCADA) systems, Plant Distributed Control Systems (DCSes), Programmable Logic Controllers (PLCs), Remote Terminal Units (RTUs), Intelligent Electronic Devices (IEDs), intelligent field devices and drives, smart meters, and other embedded industrial control and monitoring systems. These fulfill a variety of functions ranging from sensor data collection to centralized monitoring and control of entire sites or complexes of systems over large geographic areas.

Until recently, control systems operated as independent "islands of automation," providing highly reliable local control. However, over the past ten years, control networks have increasingly been integrated into the overall corporate network and are often managed by remote employees, contractors and vendors via a dialup or other connection. In this new business model, control systems become more vulnerable to online threats. Industrial environments are especially sensitive to cyber-security incidents. System downtime, loss of critical data, and loss of control over vital areas of the facility are just some of the consequences that can have a devastating impact on customers, the environment, and public safety. Even the smallest outage or performance degradation is unacceptable.

This level of risk is created by today’s trend toward using commercial off-the-shelf (COTS) technologies. Both control networks and corporate IT systems now rely on similar protocols and operating systems including Active-X, the Microsoft Windows/Linux operating system, Remote Procedure Call (RPC), Distributed Component Object Model (DCOM), Ethernet and TCP/IP. This open, standards-based functionality provides outstanding benefits, but can expose the facility to security threats that range from malicious code and attacks by hackers to operator error and technology failures.

However, this new approach, with appropriate care, also gives engineers the chance to take advantage of many of the security technologies designed to protect COTS-based systems in other industries such as finance, banking, insurance and government agencies. Many of the same security capabilities can resolve industrial problems, helping organizations to safeguard their control system networks.

The Security Challenges of Control Systems

Cyber security has not traditionally been a primary consideration in developing industrial control systems. Instead, they have generally been designed for functionality and extreme high reliability, with limited processing power and memory. Human-Machine Interfaces (HMIs) typically utilize unsecured networking stacks, common operating systems (such as DOS, Windows NT/2000 and Linux), and applications that are seldom patched after initial deployment. Many of the forms of remote access used by control networks also create vulnerabilities, including insecure communication protocols between control systems and applications such as ActiveX controls and PCAnywhere. Systems relying on these technologies may easily fall prey to viruses and worms delivered via the Internet.

However, the answer is not to simply add security technologies to industrial control networks. In fact, there is significant evidence that doing so can severely degrade control system performance. To examine this problem, the National Institute of Standards and Technology (NIST) tested typical control system processors to determine the impact of antivirus definition updates (antivirus being one of the most demanding security applications). They found that depending on the speed and loading of the processor, simply performing a virus definition update created a two- to six-minute denial of service! Other impacts may include slower operation, system shutdowns and buffer overflow, rendering hardware completely inoperable. Unfortunately, antivirus applications are only becoming more resource-intensive, and this type of performance problem continues to increase.

In addition, although today’s control networks are based on standard operating systems, they are typically customized to support control applications and may therefore be out of date. Operating system vendors frequently do not support the off-the-shelf technologies so often in use (for example, Microsoft NT4, Service Pack 2). Because of this, software patches may be either incompatible with the control system or difficult to implement without compromising service. (See Table 1 for a comparison of how control system and office IT system requirements can vary.)

Table 1: A Comparison of Control System versus Office IT Requirements for Security

Critical Differences	Control System Requirements	Office IT System Requirements
Security Skill/Awareness	Usually poor	Usually high
System Lifecycles	15-25 years	3-5 years
Patching	Slow/impossible	Frequently used
Computing Resources for Devices	Low	High
Administration	Localized	Centralized
Security Impact	Potentially life-threatening	Potentially business-threatening
Timing	Critical	Not necessarily important

For example, when the SQL Slammer worm attacked multiple control systems in 2003 via a weakness in the Windows OS, a system supplier actually sent out a letter to its customers stating that even though the worm might impact performance with a denial of service, installing the available Microsoft patch would without question shut down the control system altogether.

Another security challenge is presented by the increasing use of wireless systems. Wireless networking is significantly less costly to install than wired networks, as the cost of cabling in an industrial setting can be as high as \$3,000US per foot. However, wireless networks are vulnerable to an additional set of threats, including eavesdropping, rogue access points, interference from natural forces or reconfiguration of the physical space. With the advent of wireless-capable devices, as well as access through modems, radio and cellular links, the traditional physical security perimeter organizations once relied upon has more or less disappeared altogether.

In the midst of these challenges, the fundamental issue with securing control systems is minimizing the impact on performance. We have seen that traditional IT security policies, testing and technologies can significantly impact the operation of real-time controllers and field devices. However, low-impact IT security policies, testing and technologies do exist today and can be modified to protect the operator interfaces using commercial operating systems. For example, firewalls and intrusion detection and protection systems (IDPs) can be reconfigured to utilize control system protocols, including non-routable protocols such as DNP3 and MODBUS, and to respond to control-specific cyber incidents. They are also able to segment the myriad connections and complex architecture of the control system network into zones that permit its effective and secure management.

Threats to the Control Network

The primary sources of attacks against control systems originate via the wide area network (WAN), the Internet, and trusted third-party or remote connections. While internal threats are still significant and one of the top areas of concern for plant managers, increasing numbers of threats are originating from external sources. This mirrors the current threat trend in traditional IT systems.

Internal threats can come from a number of different sources, including attacks by disgruntled employees and contractors, or accidental infection from a device accessing the network without the latest protection and unknowingly spreading a virus, worm or other attack. However, user error and unintentional incidents actually represent the greatest risks, causing most cyber-related incidents in industrial environments. A local or remote user may access the wrong systems and make changes to them; IT personnel can perform a network penetration test that degrades performance or renders a system inoperable; or a user may download or send large files over the network and impact control traffic performance.

There is also a wide range of external threats to control systems. These range from accidental infection by a guest laptop to deliberate attacks launched by hackers, corporate spies and hostile nation-states. Today's hackers are now more often motivated by profit, with groups looking for opportunities for extortion or theft that provide a quick payoff. Such targeted intrusions are increasingly difficult to detect, which is a key reason for achieving complete visibility across the corporate network. These types of threats may include:

1. *Malicious Code (Malware)*: Malware includes the broad range of software designed to infiltrate or damage computing systems without user knowledge or consent. The most well-known forms of malware include:
 - *Viruses* that manipulate legitimate users into bypassing authentication and access control mechanisms in order to execute malicious code. Virus attacks are often untargeted and can spread rapidly between vulnerable systems and users. They damage systems and data, or decrease availability of infected systems by consuming excessive processing power or network bandwidth.

- A *worm* or self-replicating program that uses the network to send copies of itself to other nodes without any involvement from a user. Worm infections are untargeted and often create availability problems for affected systems. They may also carry a malicious code to launch a distributed attack from the infected hosts.
 - The *Trojan horse*, a type of virus in which the malicious code is hidden behind a functionality desired by the end user. Trojan horse programs circumvent confidentiality or control objectives and can be used to gain remote access to systems, gather sensitive information or damage systems and data.
2. *Distributed Denial of Service Attack*: DDOS attacks have become notorious over the past few years when used by attackers to flood network resources, such as critical servers or routers, in several major organizations with the goal of obstructing communication and decreasing the availability of critical systems. A similar attack can easily be mounted on a targeted control system, making it unusable for a critical period of time.
 3. *Rogue Devices*: In wireless networks, an unauthorized access point may be inserted into the control system to provide false or misleading data to the controller. This can cause it to issue errant commands such as triggering a failsafe device or changing operator screens to provide erroneous information.
 4. *Reconnaissance Attacks*: Reconnaissance attacks enable the first stage of the attack lifecycle by probing. This serves to provide a more focused life cycle and improve the odds of success in the attacker's favor.
 5. *Eavesdropping Attacks*: The goal of an eavesdropper is to violate the confidentiality of communications by "sniffing" packets of data on the control network or by intercepting wireless transmissions. Advanced eavesdropping attacks, also known as "Man in the Middle" (MITM) or path insertion attacks, are typically leveraged by a hacker as a follow-up to a network probe or protocol violation attack.
 6. *Collateral Damage*: This type of impact is typically unplanned or materializes as an unforeseen or unplanned side effect of techniques being used for the primary attack. An example is the impact that bulk scanning or probing traffic may have on link and bandwidth availability. Or, if a network is not properly configured, unintended traffic, such as large downloads, streaming video or penetration tests, can consume excessive bandwidth and result in unacceptable levels of network "noise" and slowed performance.
 7. *Unauthorized Access Attacks*: These are attempts to access assets that the attacker is not privileged or authorized to use. This implies that the attacker has some form of limited or unlimited control over the system.
 8. *Unauthorized Use of Assets, Resources, or Information*: In this type of attack, an asset, service or data is used by someone authorized to use that particular asset, but not in the manner being attempted.

The faster a threat can be recognized, the more quickly it can be dealt with. Preventing the behavior of the attacks and intrusions once the hacker is inside is the key to network security. There are many "back doors" and potential weak links in industrial networks. Typically, these include misconfigured devices, undocumented connections, wireless networks without proper security configurations, and open unguarded ports. A primary vector of concern is the compromise of data that can alter the operation of field devices or mislead an operator into taking inappropriate action.

Perhaps the greatest threat of all is the lack of understanding within the industrial organizations—in both Operations and IT departments—as to the seriousness of the problem. Even control system vendors still are not designing technologies for security; in fact, many are instead including vulnerable new applications such as Microsoft IIS, Bluetooth wireless communications, and wireless modems in their latest offerings.

The Impact of Security Incidents on Industrial Control Systems

The impact of cyber incidents on industrial control systems is a true daily threat, and many have already occurred. Intentional and unintentional damage to network control systems has historically included loss of confidential data, data being altered to cause erroneous equipment operation or operator information leading to mis-operation, loss of customer confidence, leaks of toxic substances and environmental disasters. In one notorious case, a pipeline rupture caused by a problem in a control network in Bellingham, WA, led to a gasoline leak into two creeks that then ignited, resulting in three fatalities, eight injuries, the shutdown of three refineries and \$45 million in damage.

Not the least of the concerns following this disaster was that, as control system networks typically have no forensics capability to help monitor and track incidents, the exact cause of the Bellingham accident could not be confirmed—though it is clearly indicative of the level of damage that can be caused by a well-targeted cyber attack. Because control systems are not designed to monitor and report network problems, many times the only indication of an incident may be the damage resulting from it. For example:

- A disgruntled contractor used a wireless link to break into a control system for Maroochy Shire Sewage, in Australia. The first 20 times the hacker entered the system, the attack was viewed as a mechanical or electrical problem with the network or its associated field devices. The perpetrator was able to intrude a total of 46 times and release millions of gallons of sewage before the problem was identified as a cyber attack. The perpetrator was later caught, but for a vehicle traffic violation, not his hacking!
- A power plant was inadvertently cycled for three hours, resulting in significant excess stress on the steam turbine rotor. It was not identified as a security incident until well after the event was over.
- In a third case, an unpatched router in a control center was attacked by a worm, resulting in the loss of communications to almost half of the utility's distribution substations. It took 24 hours just for the problem to be identified as a cyber incident.

In a laboratory test, Idaho National Laboratory (INL) and Sandia National Laboratory (SNL) technicians were able to exploit a buffer overflow by substituting compromised software within it. By using the appropriate MAC address and readily available tools, the team was able to penetrate multiple sets of security firewalls over an almost-1,000-mile area, sending compromised packets of data that took direct control of substation local area networks (LANs), as well as modifying operator screens. In other demonstrations, INL and Pacific Northwest National Laboratory (PNNL) staffers were able to change settings and create new output to modify or incapacitate operation of control centers, substations and process control valves.

While understanding the seriousness of these incidents, some plant engineers may not be quite clear on the distinction between “safety” and “security” within their systems. Because their current configuration is designed for safety, they think that it is also secure. However, while industrial equipment is typically built with several failsafe mechanisms, these are designed against circumstances that have realistic probabilities of occurring during normal operation. A cyber-security incident, on the other hand, can skew these probabilities severely, causing failures in areas that have never been considered as potential threats. For instance, fatigue failure is typically calibrated based on a normal rate of startup and shutdown. The same system could fail within a fairly short period of time, however, if the controller is intentionally or unintentionally instructed to continuously vary the rotor speed. A video released by the US Department of Homeland Security in 2007 shows a diesel power generator being destroyed as safety systems are bypassed during a cyber-based attack.

Meeting New Regulations for Network Security

In response, the Critical Infrastructure Protection (CIP) was initiated as a presidential directive by President Bill Clinton in 1998 (Presidential Decision Directive - 63) as a national program to assure the security of vulnerable and interconnected infrastructures. It was then updated on December 17, 2003, by President George W. Bush as a Homeland Security directive (Homeland Presidential Decision Directive - 7). This second directive broadened the definition of infrastructure to include physical and virtual systems that are “so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, and national public health or safety.”

The CIP program includes a national structure and a National Infrastructure Assurance Plan. It requires organizations to:

- Assess vulnerabilities to both physical and cyber attacks
- Plan to eliminate significant vulnerabilities
- Develop systems to identify and prevent attempted attacks; alert, contain, and rebuff attacks; and be able to quickly rebuild essential capabilities as needed

Today’s services infrastructures are more interlinked than ever before, as organizations increasingly rely on computer automation and Internet-based communications. Although efforts are underway, there is no unified national capability to protect today’s interrelated services infrastructure. We are still looking to understand how these relationships work, how each element functions, and how it affects the others. So far, general tests have shown that barely half of all essential services have a good security plan in place—and in some industries, such as utilities, the percentage is far lower.

In order to address today’s increased attention on network security to protect the nation’s bulk electric grid, the US Federal Energy Regulatory Commission (FERC) has approved the North American Electric Reliability Corporation (NERC) security regulations for Critical Infrastructure Protection (CIPs) for electric utilities (NERC CIP-002 through CIP-009). The NERC CIP provides the link between risk management and infrastructure assurance, providing the capability to eliminate potential vulnerabilities. With these concerns in mind, appropriately applying the new NERC requirements can help organizations to improve control system security and identify appropriate investments they need to make to safeguard their facilities both now and in the future.

A Systematic Approach to Security

One method of building security and improved reliability into control systems, without replacing the entire infrastructure, is by securing them through the network. Serving as the foundation for a variety of security functions, a properly configured network provides multiple layers of defense for every critical system and technology across the entire organization, no matter how geographically dispersed. Within this framework, facilities can be supported by a solution that blocks intrusions, detects and disables attacks, safeguards against illicit users and assures secure remote access.

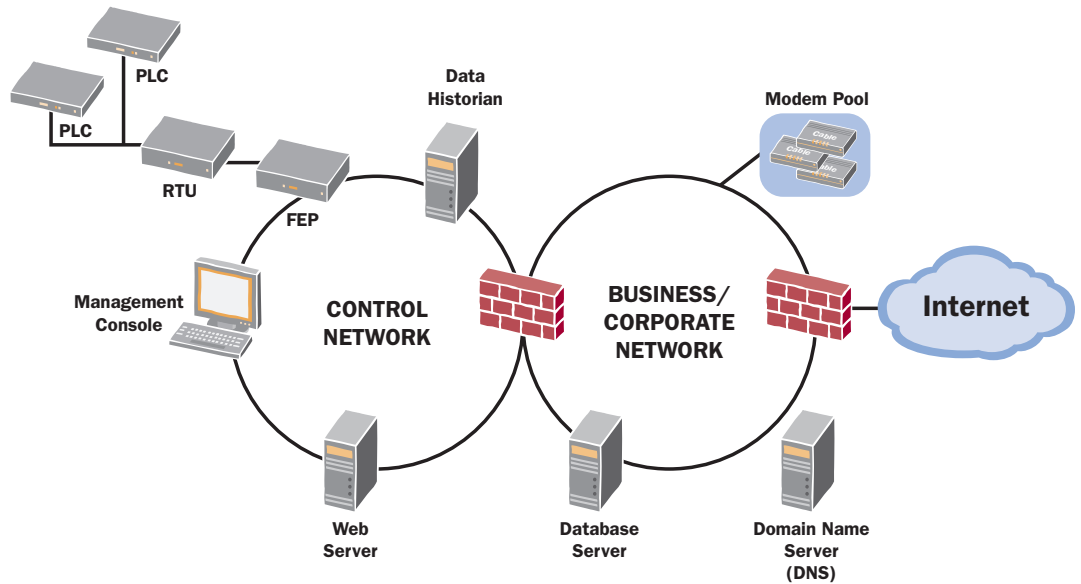


Figure 1: Today's Secured Industrial Network

A systematic approach to such a model begins by reducing the vulnerabilities of the control system network. To establish an appropriate security solution, security policies must be established for each control system to be secured with consistent rules of use, identification, authentication, and other activities defined for each level of user and role in the organization. Ideally, Operations and IT Security will work together to support this effort. This policy thus becomes the basis of security for the entire control system network, which may include servers, data historians, HMI systems, RTUs, PLCs and peripherals such as printers and network switches.

- The first step is creation of a control system security policy that details which devices, protocols and applications may run on the network, who has access to these devices from what locations, and what operations each user (or role) is allowed to perform. The Operations organization needs to identify appropriate locations to implement the policy. This can be done through configuration of controls on the devices already present on the network, and by adding various network elements to create a security perimeter and segment the network for fault containment.
- The organization must then perform a gap analysis by monitoring implementation of the policy to ensure that the new controls are effective, locate any violations, and make any corrections based on observed network behavior.
- Next, the operations team should perform a vulnerability assessment to determine whether there is any exposure that still exists.
- Finally, a risk assessment should be performed to prioritize what needs to be done and when.

Typical Security Technologies Applicable to Control System Networks

Firewalls

The first task of the industrial security solution is to avert the most obvious problems based on malicious traffic from unapproved sources. This is accomplished by the network *firewall*, a dedicated appliance that inspects all network traffic, and denies or permits access based on a pre-defined set of security policies. A firewall limits a control system's network access to specific ports and protocols from specified networks. It can also provide the ability to create distinct security zones using *Network Address Translation (NAT)*, which enables multiple areas of a private network to access the Internet using a single public IP address and *Virtual Private Networks (VPNs)*. A VPN is a communications network "tunneled" through an untrusted network (usually the Internet) to provide a segmented, secure communications path.

The firewall's main task is to regulate traffic between network segments at different trust levels—for example, between the Internet, as a zone with no trust, and the internal control network, a zone of higher trust. As a result, the control system network is not directly exposed to the public Internet or to other untrusted networks unless policies allow it. A *zone* is defined as an aggregation of network resources with similar access requirements, potential vulnerabilities, change management processes, and provision for the same consequences of security incidents. The concept of well-defined zones is especially important to ensure that appropriate policies and capabilities are applied, and security and performance requirements are met in each area. Zones also typically encompass compatible content and frequent, clearly defined communication patterns.

A third zone with an intermediate trust level, situated between the Internet (or the corporate network) and a trusted control systems network, is often referred to as a *perimeter* or *demilitarized zone (DMZ)*. The DMZ provides insulation between control areas and enterprise systems, and this is where systems and data that need to be accessed by both control and business systems and applications will reside. The DMZ can also host patch management servers, proxy servers and terminal services, enabling secure access to and from the trusted/untrusted networks.

In addition to these three zones, control system networks lend themselves to *segmentation* by function: That is, supervisory control, regulatory control and monitoring functions should be in separate zones in order to ensure isolation for fault containment. Segmenting the network enables policies that permit certain employees and contractors visibility into and control over only the systems they are authorized to access, without directly connecting to any other areas.

Today's firewalls are able to filter traffic based on source IP address, source port, destination IP address or port, protocols and domain names. However, firewalls and segmentation are only the first step in securing the network, and cannot by themselves protect control systems against unintentional mistakes or malicious attacks. An in-depth defense strategy is crucial for covering more sophisticated threats.

Intrusion Detection and Protection

A more advanced layer of defense is provided by the *Intrusion Detection and Prevention (IDP)* appliance. Such defenses (known as intrusion prevention systems, or IPS) can be deployed to help prevent attacks, or simply to detect attacks using intrusion detection systems, or IDS, also known as "sniffer" mode).

Information is sent through the network in small blocks of data known as *packets*. An IDP goes deeper than a firewall by assessing each packet based on the network protocols, the context of the communication, and its tracking of each *session* (the time the user spends communicating on the network). Akin to antivirus software on a desktop, typical IDPs contain a large repository of signatures that help to identify potential attacks by matching attempts to exploit known vulnerabilities. Such attempts might include a *buffer overflow* (where extra data is forced into the system to start overwriting memory), *protocol fuzzing* (an automated technique that sends out data packets to test for network vulnerabilities), or sending of malformed packets to crash an application. Control system-specific attacks may have signatures that differ from traditional IT attacks.

In some solutions, an IDP appliance can also provide network profile information by passively monitoring traffic and generating reports about device communications. Control system communications are generally repetitive in nature. Therefore, an IDP can baseline normal patterns of communications, and alert an operations engineer as to any deviation from it: for example, a new port beginning to send and receive data, or a new host coming online without prior notice. Such records help to establish compliance levels and provide an audit trail for forensics. This is even more useful if each session can be tied to a particular user identified as using an application. The ability to detect unknown hosts can help to identify and mitigate the threat of rogue devices entering the control network over unguarded Ethernet ports or wireless access points.

Most intrusion prevention appliances are tailored for applications in today's COTS environments, including Windows OS, Linux, Apache, ActiveX and so on. With the growing focus on critical infrastructure protection, some network vendors have integrated the ability to understand and protect protocols such as MODBUS, Inter-Control Center Communications Protocol (ICCP) and Distributed Network Protocol (DNP3) into IDP appliances. In addition to protecting critical assets, having a network device that understands these protocols enables administrators to create policies to restrict specified employees or groups of personnel to read-only access to a device, making them unable to modify operations. Such proactive policies go a long way toward preventing outages due to unintentional mistakes.

Authentication/Authorization Systems

Authentication and authorization systems protect applications by verifying user identity, providing access to devices based on that user's role and privilege level, and logging all access attempts in order to audit any infringement or misuse of critical plant functions. As most control devices such as RTUs and PLCs include little authentication or authorization capability due to lack of memory, today's security solutions provide network-based access control, which allows control systems to be accessed only by users who go through an authentication procedure that establishes a predefined identity and grants predefined privileges. The use of passwords alone is not a secure enough mechanism, yet it is still the norm to find devices in the field that rely on the manufacturer's default password. Most security standards require *two-factor authentication*, which requires the combination of two methods of identification, such as a password and a certificate. There is also a growing acceptance of *RSA key-based authentication* where the access code automatically rotates every minute, or of biometric devices such as thumb scanners.

Users are thus limited to the controls for which they are authorized, minimizing the possibility of misuse or tampering. For example, a vendor logging in to update a patch will be prevented from running any control system commands. Or, if a contractor's laptop contains any spyware or his antivirus software is not up to date, he will not be allowed access to the control network. Ideally, authorization mechanisms provide still greater control by limiting access only to certain protocols, applications and command sets.

Network Access Control

Network access control (NAC) provides the additional ability to control network and application access, based on compliance with predefined policies. This might include ensuring that users and their laptops or other devices meet a minimum baseline of security in order to gain access. Such policies can be based on various criteria, such as user identity, device identity, device health, device and/or network location. A solution including NAC ensures that the appropriate connection is properly made to the appropriate network, by both user and device. It also ensures that users and their devices meet all authentication and security policies. Since network access control applies to users as well as devices, this can become a reliable method for rogue device mitigation over wireless or wired networks.

Secure Remote Access

On top of these security capabilities, contractors, engineers and managers may remotely communicate via remote access VPNs enabled by the *Secure Socket Layer (SSL)-based security protocols*. Found in all standard web browsers, SSL provides a more secure, efficient and effective way to access control networks from an outside location or even outside the organization. This set of protocols allows secure communications via the Internet for gathering sensor data, sending instructions to field devices, performing remote maintenance and administrative data transfer tasks.

Encryption of Critical Data

Where appropriate, organizations may use *encryption* for their longer-term (non-real-time) essential control systems information. Encryption is the process of transforming information, such as a document or important message, by using an algorithm or cipher to make it unreadable to anyone who does not have the key to the cipher. It is a standard method for protecting highly confidential information. However, as heavily encrypted messages can slow network performance unless managed effectively, its use is often restricted to non-real-time messaging and data.

Monitoring for Administration and the Audit Trail

An increasingly important aspect of today's security solutions is the ability to monitor and administer the entire network to keep it at optimum performance, identify weaknesses, maintain consistent security policies, track a constant history of activity, and assure the complete safety of information. Security is a continuous process and requires diligent monitoring, reviewing and adjusting to be effective. Supported by IDP systems, firewalls, and secured switches and routers, networking solutions provide software and hardware capabilities to flag violations and serve as an audit trail. Such a solution automatically alerts engineers as to problems as well; for example, an improper connection or command will automatically trigger an alert and notify the appropriate authority immediately.

Configuration Management

A final aspect of security is helping to assure high network performance to avoid problems of availability, access and lack of service. A good security solution provides support for *configuration management* and *control*, a model that focuses on establishing and maintaining knowledge of the system and network configuration, including security. Based on this approach, operations personnel have the ability to manage security features and assurances through control of changes made to hardware, software, firmware, testing and documentation throughout the lifecycle of the systems.

The Need for Control System Security

The greatest benefit of network-based security is that it is possible to install and integrate both industrial and business systems without the expense and extensive service disruption of replacing current control system technologies. The network simply becomes another layer of functionality—one that sits on top of all the current disparate networks, data centers and control systems and unites them into a single, powerful whole. This provides end-to-end protection while providing a much greater degree of visibility into every aspect of daily operations, resulting in better security and reliability, as well as more effective management.

It is becoming increasingly important for control engineers and plant managers to meet with network security specialists to develop a better understanding of how they can best protect their vital control systems. All of the security solutions proposed in this paper are in fact already available. Such solutions must be interoperable with control system protocols in order to implement the appropriate level of security. At least one vendor has built control protocol awareness directly into its intrusion detection and prevention systems, and its solution combines IPS (to track intrusions, rogue servers and devices, as well as restrict user access to specific roles or networks, applications, or even specific MODBUS commands), firewalls, authorization/authentication and the other solutions described in this document. Ideally, the IPS is integrated with secure remote access to provide an optimal multi-layer defense mechanism.

Protection of industrial control networks is essential for maintaining a reliable infrastructure for energy generation and distribution, gas and oil drilling projects, and safe water distribution. Today we are rapidly learning that the security measures currently in place for most of these networks are inadequate. The result is unintentional power outages, growing numbers of extortion attempts, and loss of property and life. However, with a comprehensive solution in place, industrial organizations can ensure that their control systems are completely safeguarded and our national infrastructure is secure.

About the Author

Joe Weiss, Managing Partner of Applied Control Solutions (ACS) is an industry expert on control systems and security, with more than 30 years in the energy industry. Before launching ACS, he spent time at the Electric Power Research Institute (EPRI), heading programs including the Nuclear Plant Instrumentation and Diagnostics Program, the Fossil Plant Instrumentation & Controls Program, the Y2K Embedded Systems Program, and initiatives on cyber security for digital control systems. He testified to three Congressional subcommittees, and has served as a regular speaker for numerous industry events and at the NIST/NSA Security Summit.

Mr. Weiss holds two patents and has published more than 60 papers on instrumentation, controls, and diagnostics. He has also written a chapter for Electric Power Substations Engineering on “Cyber Security of Substation Control and Diagnostic Systems”. He also established and chairs the annual Control System Cyber Security Workshop, and established the International Standards Coordination Meeting on Control System Cyber Security. He has received numerous industry awards, including the EPRI Presidents Award (2002) and is an ISA Fellow and member of the ISA Engineering, Science, and Technology Policy Committee. Mr. Weiss is a registered professional engineer in the State of California and a Certified Information Security Manager.