Practical Cybersecurity for Protection and Control System Communications Networks

Scott Manson and Dwight Anderson
Schweitzer Engineering Laboratories, Inc.

This paper was presented at the 64th Annual Petroleum and Chemical Industry Technical Conference, Calgary, Canada, September 18–20, 2017, and can be accessed at: https://doi.org/10.1109/PCICON.2017.8188738.

PRACTICAL CYBERSECURITY FOR PROTECTION AND CONTROL SYSTEM COMMUNICATIONS NETWORKS

Scott Manson
Senior Member, IEEE
Schweitzer Engineering Laboratories, Inc.
2350 NE Hopkins Court
Pullman, WA 99163, USA
scott_manson@selinc.com

Dwight Anderson
Member, IEEE
Schweitzer Engineering Laboratories, Inc.
2350 NE Hopkins Court
Pullman, WA 99163, USA
dwight_anderson@selinc.com

Abstract—This paper explains practical design principles to follow for networked protection and control systems. Proven cybersecurity best practices, system architectures, monitoring methods, and defense-in-depth techniques are described. The risks and proper mitigations are explained for many common problems, such as human error, malicious malware, and advanced nation-state zero-day attacks. The paper references and summarizes several industry standards.

These insights come from the authors' design, installation, and support of dozens of operational protection and control systems. The paper is written so that a protection or control engineer with minimal network experience can easily relate to all concepts.

Index Terms—Cybersecurity, operational technology (OT), information technology (IT).

I. INTRODUCTION

The introduction of Ethernet connectivity to protection, automation, and control equipment has made practical the construction of large, sophisticated protection and control (P&C) systems. When properly implemented, Ethernet networks provide many benefits over previous communications technology, including increased process visibility, convenient remote access to engineering data, communications interoperability, and improved reliability.

However, modern Ethernet communications come with many risks. Malicious intrusions, computer operating system (OS) obsolescence, unsecured networks, uncontrolled firmware and software updates, nondeterminism, unintentional connections to outside networks, and poor network recovery can seriously harm the safe, cost-effective, and reliable operation of P&C systems.

Almost every day a cybersecurity-related intrusion is in the news. Compare this to just five years ago, when such events were almost nonexistent. This paper presents proven solutions to a number of cybersecurity challenges facing P&C system designs. There are some basic drivers challenging the security of P&C systems. Namely, there is often an assumed trust of the network, devices, and people working on a P&C system. These areas of trust are one of the fundamental targets for an attacker. Stuxnet shows that well-designed malware can inflict damage even when a P&C system is completely isolated and not attached to the Internet [1]. No one can legitimately say, "My network is isolated and therefore secure."

This paper starts with a background on pertinent industry standards and a differentiation of operational technology (OT) and information technology (IT) networking requirements. A modern P&C architecture is then described using a Purdue model. After this, a series of the most common questions about cybersecurity are answered. With these answers, the paper highlights the risks and provides recommended solutions in the secure use of Ethernet communications, OSs, controllers, and protective relays within a P&C system.

II. BACKGROUND

In the Northeast blackout of August 14, 2003, tens of thousands of people were left without power for over 24 hours [2]. The event helped push through federal regulations to investigate and resolve any potential future incidents. The directive then turned into a full regulatory system and body aimed at the security of critical infrastructure within the U.S. power grid. Today's version of the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) has evolved to support better cybersecurity and is still undergoing transformation. To its credit, and despite a number of significant issues, NERC CIP has evolved as a reference source for cybersecurity practices. NERC CIP explains what must be done but does not explain the "how to." This paper is an introduction to the "how to."

There are a number of cybersecurity standards, including NERC CIP, IEEE 1613, ISA SP99, and NIST SP800. All of these standards attempt to guide P&C system owners toward better cybersecurity. As far back as 2006, there were efforts to unify cybersecurity standards [3]; however, the use case of each industry is so different as to slow this unification.

All the modern cybersecurity standards follow the fundamental cybersecurity principles vetted during World War II—namely, policy, plans, procedures, and training. P&C system owners must determine specific risks and consequences, generate and follow a good cybersecurity policy, develop and use good cybersecurity procedures, and educate employees and contractors on those policies and procedures. In the authors' experience, a cybersecurity program is only successful when an entire organization adopts a culture of security.

III. TYPICAL P&C SYSTEM ARCHITECTURE

Fig. 1 is an example Purdue architecture diagram of a typical industrial P&C system. Both virtual and physical security perimeters must be shown in a Purdue diagram. Physical security perimeters, such as fences, locked doors, and geographic isolation, are identified at each level. Virtual security perimeters, such as firewalls and gateways are shown between appropriate levels.

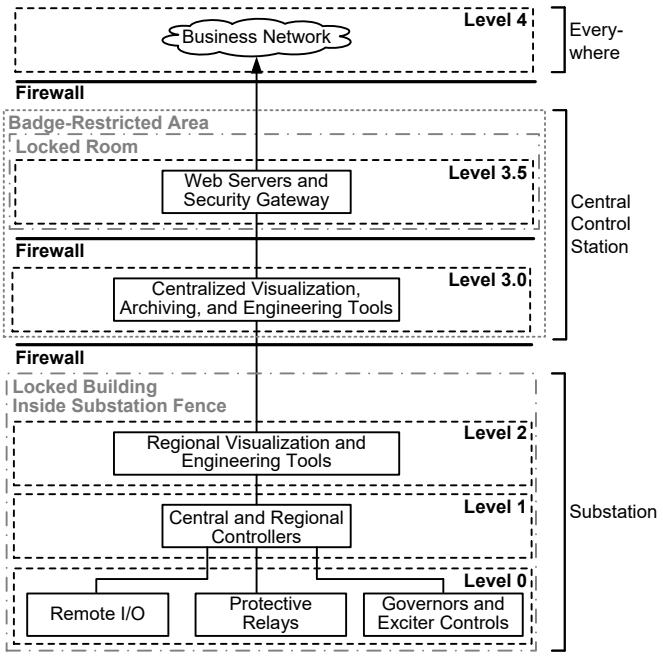

Fig. 1 Purdue Architecture of P&C System

As shown in Fig. 1, Levels 0, 1, and 2 physically exist in the electrical substation; Levels 3.0 and 3.5 exist at a central control station; and Level 4 can be a massive network spanning continents. Fig. 1 shows physical security perimeters; for example, the devices in Level 3.5 are in a locked room within a badge-restricted area; the devices in Level 3.0 are in a badge-restricted area; and Levels 0, 1, and 2 are inside a locked building inside a substation fence. Fig. 1 also shows virtual security perimeters; for example, there are separate physical firewall appliances between Levels 4 and 3.5, between Levels 3.5 and 3.0, and between Levels 3.0 and 2.

The functionality and criticality of a device determines which level it belongs to. The more critical a system, the lower it resides in the architecture. Power systems are at immediate risk of failure if Levels 0 and 1 are compromised. Most power systems can operate for weeks without Levels 2 and 3.0, and most can operate indefinitely without Levels 3.5 and 4.

Level 0 includes multifunction protective relays, on-load tap changers (OLTCs), governors, exciters, inverter controls, battery controls, and more. Level 0 protection systems minimize catastrophic damage to the most expensive assets in a power system, such as transformers, buses, lines, generators, motors, heaters, capacitors, and switchgear.

Level 1 devices are deterministic control systems that interrogate the Level 0 devices and make wide-area (multiple-substation) control decisions. Level 1 hardware devices include programmable logic controllers, front-end processors, and automation controllers. Typical P&C schemes for this level are load shedding, generation shedding, runback, localized generation load sharing, point of common coupling intertie power and power factor control and dispatching, and frequency and voltage control.

Level 2 equipment includes diagnostic and engineering tools, such as automatic event report (oscillography) retrieval, detailed sequential events record (SER) reports, and settings management for all Level 0 and 1 equipment. Human-machine interfaces (HMIs) provide the real-time status of the P&C system to operations and maintenance staff.

Level 3.0 usually includes supervisory control and data acquisition-type (SCADA-type) systems with wide-area economic controls, such as economic optimization, automated financial transactions, forecasting, and centralized visualization and historical archiving equipment. Systems must be designed such that failures at Level 3.0 have no impact on the lower, more critical levels of the P&C system.

Level 3.5 is called a demilitarized zone (DMZ). Equipment in this zone is placed in a separate network to protect Level 3.0 operational technologies from Level 4 business networks. DMZs provide visibility of the P&C system to external users, businesses, or electric utilities. Remote host sessions for visualization, such as thin clients and web servers for remote data visualization, commonly reside in the DMZ. Failure of equipment in Level 3.5 has no effect on the functionality of the lower, more critical levels.

IV. PROBLEMS AND SOLUTIONS

This section provides answers to the most commonly asked P&C system cybersecurity questions. Each problem is first explained, then best practices based on the authors' experience are shared.

A. Is Compliance Sufficient?

Security compliance is not sufficient. You must create a culture of security in the workplace and especially amongst engineering teams. Unfortunately, although well-meaning, NERC CIP has tended to create a culture of compliance rather than a strong culture of security.

To create a strong culture of security requires a concerted effort by a company's management. Building access passwords must be changed, locks and keys must be replaced regularly, employees must be continuously trained and alerted to security risks, failures must be communicated transparently, and computer access must be monitored. Data on critical equipment must be limited to a need-to-know basis, folder access must be monitored, engineering and product suppliers must be vetted, access to equipment must be limited, risk assessments must be performed on every change in a live P&C system, and more.

B. How Do I Evaluate Suppliers and Contractors?

It is critical that all persons in an organization working with control systems know and are trained on the security policy, plans, and procedures (PP&P). This is even true if a third-party contractor is hired to work on a control system. The contractor may have his own company's PP&P, but the contractor must train and work within the hiring company's PP&P. It is a common mistake for novice security planners to suggest that their security policy must be hidden or kept secure and therefore not disclosed or shared.

Equipment suppliers for critical infrastructure must be vetted to ensure proper supply chain security. Some common considerations to use in evaluating suppliers include: level of use of OS and third-party software, where the circuit boards are manufactured, where the products are assembled, where the products are designed, level of manufacturing process maturity, level of vertical integration, supply chain monitoring, quality control monitoring, and cybersecurity culture.

C. Can I Delegate Cybersecurity to IT?

It is the authors' experience that using an IT security system to protect an OT P&C system can impact the OT system's safety, determinism, and robustness. At issue with a lot of regulatory cybersecurity policies and well-meaning IT consultants is not understanding the important distinctions between IT and OT. For this reason, IT specialists usually design Levels 3.0 and higher while OT specialists design Levels 0, 1, and 2.

OT networks at Levels 0 and 1 especially require safe, deterministic, and fast-healing redundant communications [5]. For example, there is no impact on the power system if a breaker status transmission to a Level 3.0 SCADA system is delayed a few seconds. On the contrary, a power system blackout may occur if that same signal is communicated incorrectly or delayed between a Level 0 relay and a Level 1 load-shedding device. Recent advances in traveling-wave protection systems now require microsecond-level latencies in communication.

OT specialists must be experienced with the data payload, protocols, and functions occurring in the P&C equipment. They must also be experienced with the IT-related protocols used to monitor system health. For example, in a recent project, the authors designed a system that summarized and aggregated Simple Network Management Protocol (SNMP) messages retrieved from a system of 1,500 Ethernet switches (ESWs) into authenticated gateways, which converted the SNMP messages to Distributed Network Protocol (DNP3) tags. These DNP3 tags were then routed through an authenticated and encrypted tunnel to a SCADA monitoring system that is continuously monitored and maintained. This hybridized IT/OT method is an elegant and low-cost method to improve visibility of network failures without compromising security. Note that such functionality is usually not available from IT specialists because of their unfamiliarity with P&C equipment.

D. What Is My Biggest Risk?

Best practice is to perform a risk assessment to identify a prioritized list of your risks. A simplified example is shown in Table I. In P&C systems, the higher risks are often associated with high-dollar assets, such as high-voltage transmission systems, generation systems, and key buswork and transformers. Assets with long lead times or for which no spare is available are often moved to the top of a list. Be wary of single points of failure in a power system.

TABLE I
SIMPLIFIED RISK ASSESSMENT EXAMPLE

Description	Priority	Risk	Mitigation
PC-based boiler controller corrupted with malware	1	Turbines will trip offline	Replace all general-purpose OSs with embedded and whitelisted industrial controllers
Protective relay settings changed	2	Transformer destroyed after a tree falls on power line	Disable engineering access to the Ethernet port; replace Ethernet connection with serial data diode to send metering data to SCADA
Exotic transformer shot by firearm	3	Long-term outage; no spares available	Erect a concrete barricade

It is the authors' experience that the most successful (fast, efficient, and comprehensive) risk assessments are done by a multidisciplinary team composed of protection, controls, communications, operations, maintenance, and OT security experts. The most successful risk assessment teams include the personnel who designed and maintain your P&C system.

The risk assessment team must not just jump to technical controls (mitigations), e.g., firewalls, intrusion detection systems, or role-based access controls. Rushing a risk assessment can lead to unnecessary financial expenditures on cybersecurity electronics; for example, it may be a higher priority to put up a bigger fence than new routers. The team must first take the time to methodically identify risks and impacts. Only after assessing risks should a team choose a security model to protect a P&C system.

The risk assessment team must keep in mind that priorities are different at each level. For example, for Levels 0, 1, and 2, you must give highest priority to network availability, followed by data integrity, and lastly data confidentiality (AIC). For Levels 3.0 and up, you must give highest priority to data confidentiality, followed by data integrity, and then lastly data availability (CIA). IT personnel learn their trade in an environment of CIA priorities, whereas OT personnel learn their trade in an environment of AIC priorities.

E. Can I Merge My Process and Electrical Networks?

Don't do it. Large integrated networks offer a simpler target for denial of service attacks. Large networks are much harder to debug, network outage times generally increase, and, hence, reliability decreases. The convergence times of automatic network reconfiguration, such as Rapid Spanning Tree Protocol (RSTP), grow as networks get larger. The authors strongly advise dividing up large networks into smaller sections with OT protocol gateways at the boundaries.

In one dramatic example of the dangers of large networks, the authors successfully brought down best-in-class IT systems by injecting a small amount of normal P&C data traffic and then turning on and off network interface card (NIC) ports to cause RSTP reconfiguration. Fig. 2 shows how chipsets used in IT ESWs create a message storm when faced with high volumes of incoming data. If the IT ESW detects messages from unknown Internet Protocol (IP) addresses, it sends the messages to all ports. This means that during the time of RSTP recovery, the IT ESW will itself cause a storm of data. This causes data losses on P&C networks that use commercial-grade IT Layer 3 ESWs; this in turn can create power system blackouts. The problem will occur whenever RSTP media access control (MAC) address tables within the IT ESW are flushed and relearned. This problem does not exist with software-defined networking (SDN) switches designed for the OT environment.

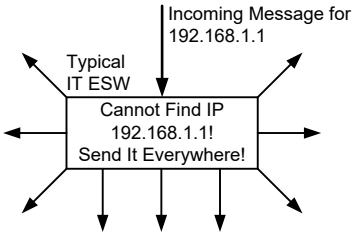

Fig. 2 IT Switches Do Not Belong in OT Networks

Most dangerous is mixing Layer 2 protocols like IEC 61850 with IT switches. Layer 2 protocols cannot be routed and exit every port of a switch unless virtual local-area-network-type (VLAN-type) methods are used. Even small network storms caused by Layer 2 protocols are well known to shut down critical Layer 0 protective relays from some suppliers [6].

F. How Do I Document What I Have?

To identify cybersecurity vulnerabilities, it is essential to identify and classify all data moving around a network [7]. Proper documentation of a P&C system requires three essential types of architecture diagrams:

1. Purdue diagrams, such as Fig. 1. These represent a simplified network topology/architecture that allows engineers to communicate their design to IT and management. Networks with more than 10,000 intelligent electronic devices (IEDs) fit into a single C-size Purdue drawing. Purdue-type diagrams are commonly created in hours.
2. Physical architecture diagrams, as shown in Fig. 3. These show 100 percent of the equipment and physical port connections, switches, relays, NICs, media converters, computers, and physical media (fiber optics [FO], copper, or radio). Everything in this document must be physical and not virtual. It may take many drawing sheets to show all the connections. Physical architecture diagrams often take weeks to create.
3. Data flow diagrams (DFDs), as shown in Fig. 4 (which is the DFD of Fig. 3). DFDs show 100 percent of the communications on a network. A DFD is the most important P&C system cybersecurity-related document and also the most commonly omitted. No physical media or cabling are shown on a DFD. Every single data communication is identified by direction, protocol, and a summary of the data content. Software and/or firmware components on the producers and consumers are identified. A DFD is the most essential document for the programming, debugging, and maintenance of a large P&C system. DFDs usually take many drawing sheets to show all of the network traffic and often take months to perfect.

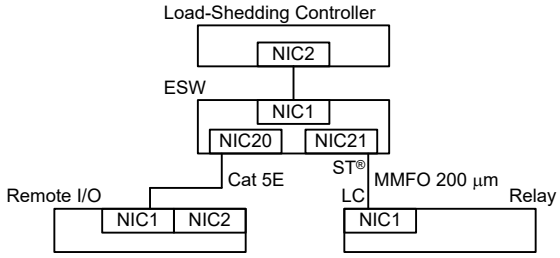

Fig. 3 Example Physical Architecture

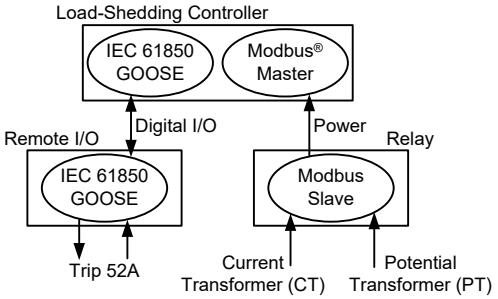

Fig. 4 Example DFD for Fig. 3

G. How Do I Know if There Is Malicious Traffic?

DFDs are commonly validated by monitoring network traffic. The traffic on a new P&C system is monitored prior to delivery. This monitored traffic can be used to create a baseline by identifying normal network traffic. Any traffic not recognized must be chased down and either accepted or corrected.

Many software "sniffing" tools are available to capture and decode Ethernet traffic. Sniffing software commonly runs on a laptop, and the setup is trivial. However, locating the best place to monitor the traffic takes some understanding of the network architecture. Much like the observer effect in physics, the act of capturing the data can commonly invalidate the

data. To avoid interfering with the data being collected, the most common solution is to set up an ESW port as a mirrored port that echoes all the traffic of another port of your choice.

For example, in Fig. 5, the ESW is configured to mirror all of the data received and transmitted on NIC1 to NIC2. This allows the laptop to see all the Modbus and IEC 61850 Generic Object-Oriented Substation Event (GOOSE) traffic going to and from the load-shedding controller of Fig. 3 and Fig. 4. Note that SDN-based mirroring is also used for this purpose because of its inherent ability to route and duplicate traffic.

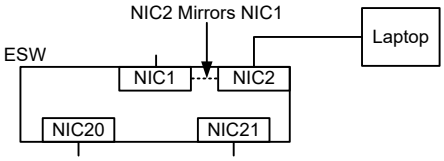

Fig. 5 Mirroring the Load-Shedding Data From Fig. 3

Software designed for monitoring and alarming on malformed or unexpected traffic [8] is also available. These alarms should be transported to Level 3.0 equipment so that operations personnel can call in an OT investigation.

H. Has Someone Connected an Outside Network?

Under no circumstances should a P&C system ever connect directly to the Internet. The P&C system gains some security from mistaken Internet connections by using private IP addresses. If public IP addresses are used in a P&C system and someone makes a connection to the Internet, the P&C IP will be accessible to anyone on the Internet. If a company does not own sufficient IP addresses, there are a large number of reserved IP addresses that can be used in a P&C system [9].

There are several software tools used to scan computers on a network for known vulnerabilities [10]. These tools raise an alert if any vulnerabilities are detected that malicious hackers could use to gain access to any computer connected to a network. There are also software tools that actively look for connections to the Internet or World Wide Web [11]. It is best practice to run these scans as part of an OT factory acceptance test (FAT) and prior to accepting new equipment onsite. This system should be used carefully on a live (production) P&C system.

I. How Do I Know if the Network Is Functioning?

Switches, firewalls, or software on a computer cannot know for certain if the network is functioning adequately on a P&C system. Only the P&C equipment sending and receiving data packets can make this determination. This is because there can be hundreds or thousands of firmware stacks, NICs, microprocessors, and software modules involved in a single transaction of data. For example, a communications driver error in a relay will not be detected by a switch.

As a final check to confirm that all communications are operating acceptably, data producers are commonly set up to send a cycling “watchdog” signal to a remote data consumer, as shown in Fig. 6. Note how the remote I/O module in Fig. 6 toggles a bit inside a logic processing environment running inside a central processing unit (CPU). This same CPU sends the communicated message through a driver, which then transports the message to a NIC. The NIC contains firmware that buffers and sends the message out on hardware in the NIC.

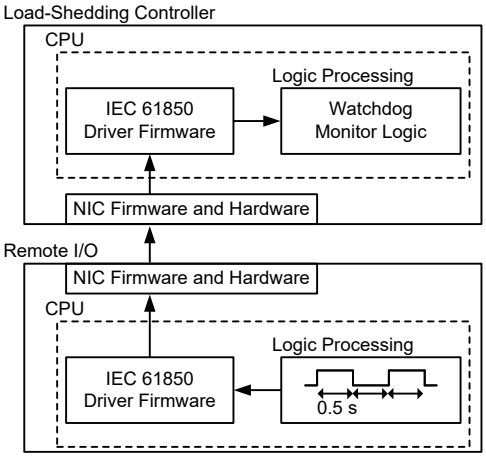

Fig. 6 Watchdog Monitoring Example

In Fig. 6, if the data consumer detects that the signal has stopped changing state or has unacceptable latencies or indeterminism, then the logic processing inside the load-shedding controller alarms the channel as failed.

Only the methods shown in Fig. 6 have proven to be 100 percent successful in properly assessing the health of P&C communications systems. The data provided from the watchdog method are commonly captured into statistics on “channel availability measurements,” which are annunciated to a Level 3.0 engineering workstation.

Once a communications channel has been detected as failed, other tools are employed to further deduce the root cause of the problem. Engineers can use SNMP, log files in controllers, and Sequence of Events (SOE) records in Level 0 and 1 equipment to further isolate the cause of the communications failures.

J. How Do I Handle an Obsolete OS?

Commercial OSs become obsolete every few years. Security patches cannot be gained and new software will commonly not function for an obsolete OS. For these reasons, the authors advise P&C system designers to only use general-purpose OSs on devices at Levels 3.0 and up. Commercial OSs should not be allowed into lower levels of a P&C system.

The authors advise replacing computers running general-purpose OSs with embedded controllers whenever possible. The authors have successfully replaced many personal computer (PC) control systems with embedded controllers.

In situations where an embedded controller cannot be used to replace a PC, the user has only two choices: 1) replacing the PC, OS, and software or 2) whitelisting.

Upgrading an obsolete OS is commonly a very expensive and inconvenient measure. This is because the software

running on the obsolete OS needs to be completely revised, replaced, reconfigured, and retested. This process then needs to be repeated every few years as the OS is again obsoleted. The next section discusses whitelisting.

K. Is Whitelisting Practical?

Whitelisting is a method that allows only known software components to run in memory [12]. Whereas antivirus methods attempt to protect the computer by looking for known viral signatures, whitelisting stops all unknown processes from running in memory.

The authors have significant practical experience with whitelisting both general-purpose OSs and embedded controllers. Whitelisting on a general-purpose OS can prove challenging because of the large number of software components used on a general-purpose OS. On the other hand, whitelisting on an embedded controller or protective relay is a very successful method.

Whitelisting on a general-purpose OS is viable in the authors' experience when it is thoroughly designed and tested in a simulated environment. The design, implementation, and testing must occur long before an OS goes into use on a P&C network. Trying to make whitelisting work reliably on a full-sized, general-purpose OS has proven impractical, hence all such OS whitelisting requires a significantly reduced OS size.

Whitelisting on embedded devices, such as port servers, security gateways, controllers, communications processors, and protective relays, is strongly recommended. Embedded devices typically are application-specific and do their job and nothing else (unlike a general-purpose OS). Whitelisting on an embedded device protects the kernel, applications, and memory access.

L. What Is a Lifecycle Strategy?

IT staff may ask for a "lifecycle management refresh strategy." The best strategy to manage a P&C system is to design a Level 0, 1, and 2 system without OSs, with high-availability equipment, and with equipment from a supplier that openly discloses vulnerabilities, patches, and firmware releases.

The cost of replacing obsolete equipment quickly outweighs the additional costs of delivering a system composed of hardened embedded components. IT experts are often surprised when they discover that most P&C systems are designed for 25 years of service. Standard commercial-grade IT equipment has a design life of 3 years.

M. How Do I Protect the System From Disgruntled Former Employees?

Twenty years ago, this involved changing the locks. Today, this additionally requires password changes, deletion of user privileges, and more. This can be done with user authentication software running on a server [13]. These credentials can be a custom setup to allow individuals access to a limited set of equipment. For example, this method is used by IT to limit employee access to only their own computers. These credentials are managed by domain controllers that manage a group policy distribution. A group policy is basically a database of privileges given to users.

Fig. 7 depicts an example of how a person logs into a security gateway (also known as an authentication proxy server) to gain access to a relay. In this process, the security gateway communicates with a domain controller to validate the access rights of a user. The security gateway is in charge of randomizing the passwords on the relays; thus, the proxy cannot be bypassed.

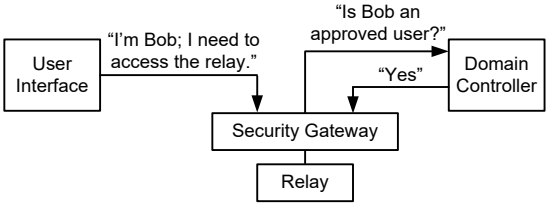

Fig. 7 Remote User Authentication

The system shown in Fig. 7 can also be configured to support the use of Remote Authentication Dial-In User Service (RADIUS). RADIUS is a form of multifactor authentication. An example of multifactor authentication is the constantly changing passkey that some people carry around on their keychains.

N. How Can I Minimize Firewall and Router Costs?

Large numbers of firewalls and routers represent a significant maintenance cost. Their firmware is constantly updated to adapt to the latest attack. Therefore, the question morphs to "How do I inexpensively and safely send data to another system without firewalls and routers?"

There are alternatives to Ethernet and firewalls at Levels 0, 1, and 2 that provide similar performance without the dangers of using IT-based Ethernet protocols. One alternative communications backbone is constructed with serial-based multiplexed technology. The data flowing on these channels are segregated into real- and non-real-time channels to ensure deterministic and prompt delivery of status and controls data. An example of a multiplexed network is shown in Fig. 8.

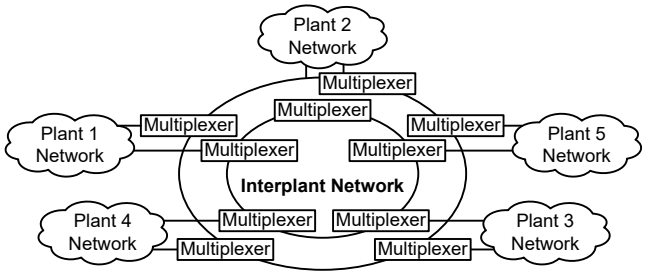

Fig. 8 Serial Multiplexed Alternative to Ethernet [5]

One inexpensive and effective alternative to a firewall is a serial data diode technique, as depicted in Fig. 9. This method uses a single directional (unsolicited) serial data channel between devices at different levels in the Purdue diagram (refer to Fig. 1). The transmitter gateway is wired to the receiver gateway with only a transmit hardware line (two

wires—transmit and ground). Without a receive line, there can be no exchange of data back into the P&C system. Universal asynchronous receiver/transmitter (UART) integrated circuits used only as transmitters do not respond to remote controls when their receivers have no hardwired line connected. This design has no way for the P&C system to receive data from the outside world.

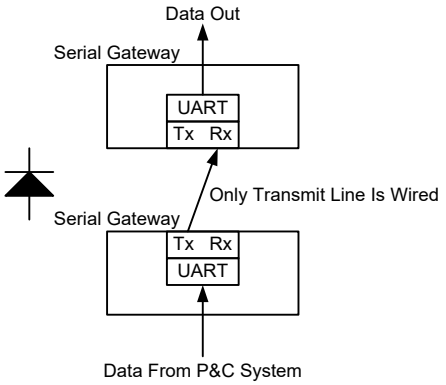

Fig. 9 Serial Data Diode

There are some simple techniques to provide defense-in-depth security without routers and firewalls. These methods to prevent outside intrusion focus on diverse tactics to dissuade all but the most skilled intruders. For example, unused communications ports are shut down. Every used port on ESWs is configured to communicate only with a particular MAC address. Data flows configured within SDN prevent any outside data traffic from reaching critical P&C equipment [14]. Obscure FO connectors (such as ST and LC), FO sizes uncommon in the IT world (such as 200-micron core), and light wavelengths not commonly used by IT can add low-cost security. All physical communications ports must be behind locked panel doors.

The authors only suggest the use of routers in the Level 3.5 DMZ. The authors discovered that the bandwidth limitations of Ethernet routers are cause for concern at Layer 2 and below. On a project protecting the western interconnect of the U.S., best-in-class routers were causing millisecond-level delays in package delivery [15]. For this reason, routers were designed out of the project and replaced with deterministic protocol gateways that provide guaranteed response times. It is a common experience that IT routers are slow for package routing, and this is a critical reason why the authors only use SDN and protocol gateways for critical deterministic communications. SDN and protocol gateways are proven to ensure repeatable and acceptable latencies and higher reliability in message delivery.

O. How Do I Route Traffic Flow?

Ethernet networks with rings have alternative routes over which messages can flow between devices, as shown in Fig. 10. The ESWs autonegotiate which route is the primary and backup for normal data flow based on user settings. The most common protocol used by switches to autonegotiate the flow of traffic is RSTP, although many other protocols are available.

It is the authors' experience that large redundant P&C system networks with multiple rings (or ladders) can constitute a significant network design challenge. RSTP recovery times can grow unacceptably large if the ESW settings are improperly designed or tested. Hire an OT expert who understands the ESW to design these large systems. Test every failover mode during an OT FAT.

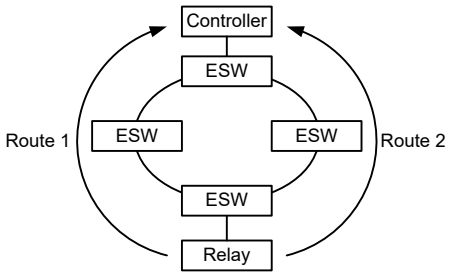

Fig. 10 Typical Ethernet Ring Network

When possible, P&C systems should use SDN. SDN allows all communications paths shown in a DFD to be hard-coded into all ESWs on a network, thus denying any outside traffic from operating on the network [11].

P. How Do I Rebuild a Failed OS?

It is the authors' experience that installing software and retesting a Level 3.0 application can be tedious, be time-consuming, and require a 100 percent retest of functionality. This is unacceptable in a live plant situation. For example, rebuilding a simple HMI running on a commercial OS can require reinstalling and relicensing an OS, reinstalling several software packages, and then loading the HMI configuration. This is commonly complicated by OS and software obsolescence and incompatibility.

Virtual machine (VM) software or binary backups are strongly encouraged. A faster system backup and recovery plan is to use VMs or binary image backups. VM software acts as a vessel for installation and normal operation of an OS. VM software, for example, can allow many different OSs to operate seamlessly on the same computer. Binary images are stored copies of every bit on a computer hard drive. These binary images can be restored to a computer with boot-level software. VM software commonly allows a failed computer image to be restored in minutes. Binary images take longer to restore, but they are commonly lower-cost solutions.

Note that the correct and easiest solution is to avoid general-purpose OSs altogether.

Q. What Are My Alternatives to General-Purpose OSs?

Level 0 and 1 equipment must *not* contain a general-purpose OS. It is necessary to limit the use of general-purpose OSs because of the large risk and cost they cause to P&C systems. General-purpose OSs are used primarily because of their ability to run data archiving, visualization, and engineering software tools. One alternative to an HMI running on an OS at Level 2 is to use embedded web servers running in Level 1 embedded controllers. Another is to only retain the

historical data the Level 0 and 1 equipment can retain on their own, which can be significant.

R. How Do I Log Behaviors?

In P&C systems, there is a distinction between logs that occur from the functional P&C system and logs that occur from OT equipment. The authors recommend maintaining both types of logging data for P&C systems.

For example, an SOE log exists in most Level 0, 1, and 2 P&C devices (such as controllers and relays) to monitor equipment state changes. The SOE logs help an engineer diagnose a problem with the protection or control system.

In contrast, OT event logs commonly capture failed user login attempts or communications events with SNMP traps. Syslog is a common IT protocol used for sharing logs between OSs at Levels 3.0 and above. There is a variety of syslog collection and viewing software available.

It is common to store both syslog and SOE records on a network-attached storage (NAS) array for years of retention. The long-term storage of data is essential for capturing long-term metrics on P&C system performance.

S. How Do I Manage Updates?

Antivirus, software, OS, firmware, or device configuration settings may require updates. All updates must be tested in a realistic testing environment before being put into a production system. This testing environment must contain a representative set of the P&C equipment, which must be fully configured and operational during testing. For example, updates to the firmware of an ESW must be tested on a live (offline) system with active watchdog communications monitoring. Users are strongly advised to read and understand supplier service bulletins and perform risk assessments before undergoing updates.

T. How Should Suppliers Protect Our Information?

Protecting sensitive information requires both physical and virtual controls. An example of physical controls is a lock on a door, whereas an example of virtual controls is the password to a computer. Consulting, engineering, and manufacturing suppliers must have incident response procedures in place in case they have a breach of data.

Access to customer information must be limited to personnel listed in the project execution organizational chart. This organizational chart is controlled as confidential. Access privileges are assigned to all personnel on the chart. All documents should then be given a designation of either “Public,” “Confidential,” or “Classified.”

Training is essential for the success of a cybersecurity program. Supplier employees should attend company training for security, business processes, and safety as required to ensure incident-free implementation of P&C systems.

Supplier risk management includes the transparent communication of risks from supplier to user. Service bulletins must be supplied by the supplier to inform the user of risks of misoperation, loss of data, or possible outside intrusion because of defects found in a product.

Background screening, training, and regular employee monitoring must be done by all suppliers. They must track every component and setting in their system. Suppliers must keep traceable records of all designs, documents, and manufactured components back to a supplier, a specification, and a test sequence.

U. How Do I Get Secure Remote Access?

Remote access to a P&C system can mean two different things, each of which has different design requirements. The first is allowing an engineer remote access into the P&C system for maintenance or disaster recovery with sometimes significant access privileges to Level 0, 1, and 2 devices. The second is transportation of P&C data to levels above the DMZ for monitoring. The second implementation may allow some limited controls, such as alarm acknowledgement.

For remote engineering access, the authors recommend using an encrypted and authenticated host-to-host connection from the outside world into the DMZ. This is accomplished by placing a server in the DMZ that acts as an intermediate access point, commonly called a “jump server.”

Monitoring data and controls are commonly transported out of a DMZ to control centers manned with operators. If only monitoring is required, this form of remote access is best accomplished with serial data diode techniques or thin client applications that reside on a DMZ data server.

V. CONCLUSIONS

Key conclusions of the paper are as follows:

1. Always use defense-in-depth methods, including diverse methods of both physical and virtual security.
2. Security compliance is not sufficient. A cybersecurity program is only successful when an entire organization adopts a culture of security.
3. Use SDN and protocol gateways to break up large networks into multiple smaller, deterministic networks.
4. Perform a detailed risk assessment with a multidisciplinary engineering team before spending money on cybersecurity.
5. A DFD is an essential document used to identify and classify all data moving around a network.
6. Always use watchdogs embedded into devices.
7. Large numbers of firewalls represent a significant maintenance cost. Serial data diodes are a significantly lower-cost and simpler alternative to firewalls.
8. Minimize your use of commercial IT equipment.
9. Only select suppliers and engineering services teams that protect your information.
10. Perform a comprehensive OT FAT before accepting equipment from a supplier.

VI. REFERENCES

[1] D. Kushner, “The Real Story of Stuxnet: How Kaspersky Lab Tracked Down the Malware That Stymied Iran’s Nuclear-Fuel Enrichment Program,” February 2013. Available: http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet/.

[2] U.S.-Canada Power System Outage Task Force, “Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations,” April 2004. Available: http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf.

[3] B. McNaught and M. Hadley, “Control Systems Security Program (CSSP),” presented at the 2006 International Control Systems Security and Standards Coordination Workshop, Portland, OR, August 2006.

[4] D. Shinder, “SolutionBase: Strengthen Network Defenses by Using a DMZ,” TechRepublic, June 2005. Available: http://www.techrepublic.com/article/solutionbase-strengthen-network-defenses-by-using-a-dmz/.

[5] K. Ravikumar, T. Alghamdi, J. Bugshan, S. Manson, and S. Raghupathula, “Complete Power Management System for an Industrial Refinery,” proceedings of the 62nd Annual Petroleum and Chemical Industry Technical Conference, Houston, TX, October 2015.

[6] N. Seeley and K. Concienne, “Making Peace With Communications Networks: What Power Engineers Need to Know About Modern and Future Network Communication for Plants and Substations,” proceedings of the 58th Annual Petroleum and Chemical Industry Conference, Toronto, Canada, September 2011.

[7] R. Bradetich, “Framework for Evaluating Information Flow Security in Multicore Processors,” University of Idaho, Moscow, ID, 2012.

[8] K. Scarfone and P. Mell, “Guide to Intrusion Detection and Prevention Systems (IDPS),” Computer Security Resource Center (National Institute of Standards and Technology), SP800-94, February 2007. Available: http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf.

[9] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, “Address Allocation for Private Internets,” Network Working Group, Best Current Practice, February 1996. Available: https://tools.ietf.org/html/rfc1918.

[10] The Government of the Hong Kong Special Administrative Region, “An Overview of Vulnerability Scanners,” February 2008. Available: http://www.infosec.gov.hk/english/technical/files/vulnerability.pdf.

[11] D. Goldman “The Internet's Most Dangerous Sites,” May 2013. Available: http://money.cnn.com/gallery/technology/security/2013/05/01/shodan-most-dangerous-internet-searches/.

[12] D. Shackleford, “Application Whitelisting: Enhancing Host Security,” SANS Institute Reading Room, Bethesda, MD, October 2009. Available: https://www.sans.org/reading-room/whitepapers/analyst/application-whitelisting-enhancing-host-security-34820.

[13] “Active Directory,” Microsoft TechNet, January 2005. Available: https://technet.microsoft.com/en-us/library/cc782657(v=ws.10).aspx.

[14] Open Networking Foundation, “Software-Defined Networking (SDN) Definition,” *SDN Resources*. Available: https://www.opennetworking.org/sdn-resources/sdn-definition.

[15] S. Manson, D. Miller, R. Schloss, S. Raghupathula, and T. Maier, “PacifiCorp’s Jim Bridger RAS: A Dual Triple Modular Redundant Case Study,” proceedings of the 11th Annual Western Power Delivery Automation Conference, Spokane, WA, April 2009.

VII. VITAE

Scott Manson received his M.S.E.E. in electrical engineering from the University of Wisconsin–Madison and his B.S.E.E. in electrical engineering from Washington State University. Scott is currently the Engineering Services Technology Director at Schweitzer Engineering Laboratories, Inc. In this role, he provides consulting services on control and protection systems worldwide. He has significant experience in power system protection, real-time modeling, power management and microgrid control systems, remedial action schemes, turbine control, multi-axis motion control, web line control, robotic assembly, and precision machine tools. Scott is a registered professional engineer in Washington, Alaska, North Dakota, Idaho, and Louisiana.

Dwight Anderson received his B.S. in electrical engineering from Steven’s Institute of Technology. He is now a security engineer for Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington. Prior to joining SEL in 2005, he worked 20 years for Hewlett-Packard as a business development manager and systems engineer, working on projects ranging from signal intelligence systems to SCADA system programming. He is an active member of the FBI InfraGard team. He is a registered professional engineer in Texas and a Certified Information Systems Security Professional (CISSP).

Previously presented at the 64th Annual Petroleum and Chemical Industry Technical Conference, Calgary, Canada, September 2017.

20170328 • TP6745