In the News: Recent Security Failures Prompt Review of Secure Computing Practices

David Dolezilek, Bryan MacDonald, Jason Kraft, and Peter Dolezilek
Schweitzer Engineering Laboratories, Inc.

Presented at the
Power and Energy Automation Conference
Spokane, Washington
March 26–28, 2013

Originally published in the
proceedings of the 2nd Annual Protection, Automation
and Control World Conference, June 2011

In the News: Recent Security Failures Prompt Review of Secure Computing Practices

David Dolezilek, Bryan MacDonald, Jason Kraft, and Peter Dolezilek, *Schweitzer Engineering Laboratories, Inc.*

Abstract—**Recent unintended events involving control system computers have contributed to injury and death, ecological catastrophes, and undetected subversion of applications. Causes include critical hardware failures, lack of redundancy, disabled failure alarms, virus infection of islanded computers, and undetected automatic execution of unwanted programs.**

Control system best engineering practices call for the use of a purpose-built computing platform with embedded operating system and applications whenever possible. However, many contemporary control systems include general purpose computers with a Linux® or Microsoft® Windows® operating system for the convenience of using third-party software applications. Unfortunately, general purpose computer features do not discriminate between intended and unintended applications, may be needed only during initial configuration, or are completely unwanted but automatically included. These well-known features, such as startup routines, power-off switches, and USB, serial, and Ethernet ports, become unintended security flaws when left unattended. Further, the construction of general purpose computers becomes a security liability when devices with low mean time between failures (MTBF) fail or restart without an alarm indication and disrupt or suspend critical applications.

Detected and undetected threats come from hackers, disgruntled employees, terrorists, and countries with sophisticated information warfare plans and capabilities. Leadership to create and manage a cybersecurity program involves every aspect of company personnel and processes for compliance. However, recent events illustrate how engineers and technicians who design, build, and operate these control systems are truly the first line of defense against a security breach.

Though most security failures remain unreported or private, this paper uses recent public security failures to illustrate the need for, and use of, secure computing practices within the design and safe operation of control systems, including the following:

- **Verify (do not assume anything).**
- **Use fault-tolerant hardware and software applications.**
- **Design control system quarantine failures.**
- **Secure the platform, and disable extraneous features.**
- **Protect against malware.**
- **Restrict entry points, such as removable media (DVDs), USB, and network connections.**
- **Block malicious software with a firewall.**
- **Back up data.**
- **Use strong passwords.**
- **Verify third-party software.**
- **Maintain a network access control list (ACL).**
- **Test and apply patches.**
- **Install physical security.**
- **Review logs and alarms.**

This paper provides a road map for both compliance and day-to-day operations to maintain safe and secure control systems for any industry.

I. Introduction

If you read the news, you know that control systems and critical infrastructure are increasingly becoming targets and victims of cybercrime. Last year, news broke about Stuxnet, an incredibly sophisticated worm designed to sabotage control systems. While this most sophisticated, targeted worm broke through layers of defense to damage specific power systems, the vast majority of failures require far less sophistication. In fact, a coordinated attack labeled Night Dragon targeted efforts by cybercriminals to steal information from several large energy companies using common tools [1]. Regardless of the level of sophistication, one thing most attacks have in common is an expectation that human defenses will fail. In these industry-specific attacks, attackers count on humans to fail.

There are numerous documents available to help design and apply an appropriate company and system security posture, such as “Ten Tips for Improving the Security of Your Assets” [2]. However, this paper addresses the responsibility of the personnel operating computers in a control system or substation automation system. Because computers provide more and more useful information to operators, they become larger targets for unauthorized access. They also become more critical to the ability of operators to perform their duties effectively. This paper discusses practical methods to know the status of your system and computers, protect them from unauthorized access, and choose the best replacement to support business continuity or disaster recovery.

Operators are responsible for the health and security of the operator workstation, which is the human-machine interface (HMI) to the system. This paper itemizes best practices for people to follow to protect control system operator computers from security failures.

Depending on the industry, the impact of downtime can be disruptive, destructive, or deadly. The use of highly available computers, measured as a high mean time between failures (MTBF), is as important as designing a highly available process control system. Dangerous and expensive computer downtime costs are a real threat to many industries, including power and energy, pharmaceuticals, oil and gas, water, and wastewater.

One business impact of downtime is dollars lost per hour because of unavailable systems. A typical computer may take 2 hours to repair; however, if not maintained while in service, the replacement after failure can take 16 to 24 hours. Typical annual costs of downtime for these systems vary from $5,000 to $50,000. The financial return on investment of a rugged, high-availability computer can be as short as one month. The other impacts of choosing hardened, high-availability computers include better operator safety, peace of mind, and success.

II. Select Hardened Computers for Business Continuity and Disaster Recovery Replacements

Because hardware reliability is such a crucial element to system success, well-established and standardized analysis techniques and methods are used to analyze and measure device reliability to identify and remove areas of weakness. The field of reliability engineering is devoted to the development of tools and techniques for use in choosing appropriate system design and devices, including computers used for operator workstations.

The primary responsibility of the operator, with respect to the operations computer, is to physically secure the device and observe and react to logs and alarms. However, in some instances, the operator is involved in planning for or performing the replacement of control system hardware. This replacement of devices maintains business continuity when computers fail because of insufficient ruggedness for the environment or during disaster recovery replacement of a computer damaged in a control system or process accident.

There are many metrics commonly used in reliability engineering to help assess device reliability. These metrics are generally used to make topology decisions and device selections during system design. However, they are also used to make device replacement selections, such as operations computers, that match or improve the system reliability baseline.

III. Available Reliability Analysis Techniques

Because the reliability field is broad, there are many different types of reliability analysis techniques available. These techniques include reliability evaluation, risk assessment, and topology analysis tools. Device reliability is directly related to the percentage of time that the computer is in service and available to do its job. Therefore, device reliability assessment is sufficient for choosing fault-tolerant replacement hardware. The other techniques are used during the initial control system design.

The following parameters describe some of the most common reliability measures that can be obtained from reliability analyses. While there are many more reliability measurements available, the following parameters are widely used:

- Failure – a device has failed any time that it does not perform as specified or, in the absence of a specification, deviates from reasonable expectation of performance.
- Failure rate – the expected rate of occurrence of failure or the number of failures in a specified time period. Failure rate is typically expressed in failures per year. For example, you would predict a single failure during a five-year period if your computer has a failure rate of once every five years. Or, if you are using five computers simultaneously, you could predict one or more failures each year.
- MTBF – actual incidence of field failures for a large population of installed units. MTBF is the accumulated number of years in service of products divided by the number of products returned because of failure. MTBF is the inverse of the failure rate. This means that if, for example, 1,000 computers from a manufacturer are in service for a year and 4 computers experience failure, this computer has an observed MTBF of 250 years (1,000 service years/4 failures).
- Reliability – the probability that the item will perform a required function without failure under stated conditions for a stated period of time. Reliability is significant because it takes into account time. The measure of reliability answers the question: “How likely is it that my system will remain operational over a period of time?” Because reliability is expressed as a probability, it is always a value between zero and one.
- Availability – also a probability value, availability indicates the probability that a system is operating at a particular point in time. It answers the question: “How likely is it that my system is operating at *x* hours?” Availability differs from reliability because it factors repairs into the measurement. To determine availability, the time to perform a repair must be known. Because availability is expressed as a probability, it is a value between zero and one.
- Mean time to repair (MTTR) – average time to return a failed item to an operable state. MTTR is normally expressed in hours and indicates how long it takes to repair a system that is down due to a failure. Generally, MTTR does include logistics time, such as the time required to receive a replacement part. However, for this specification, MTTR indicates the actual time it takes to correct the problem once on location with tools and replacement components.
- Unavailability – the complement of availability. It is a probability value between zero and one that indicates the likelihood that a system or device is not operational at a specified point in time.

IV. Reliability Analysis Provides Better Systems and Better Performance

The importance of reliability for device and system success is undeniable. To accurately track, measure, and improve reliability parameters, a wide array of techniques have been developed by device manufacturers and system designers. In fact, IEC 61850-3 standardizes reliability and maintainability metrics [3].

IEC 61850-3 makes frequent reference to IEC 60870-4, which specifies performance requirements for a telecontrol system, classifying these requirements according to properties that influence the performance of the system [4]. IEC 61850-3 Section 4 describes internationally standardized requirements for the quality of substation communications systems and has the following scope:

> [It] details the quality requirements such as reliability, availability, maintainability, security, data integrity, and others that apply to the communications systems that are used for monitoring, configuration, and control of processes within the substation. [3]

The standard goes on to say that each device in the system, including computers, should be chosen considering the graceful degradation principle from IEC 60870:

> There should be no single point of failure that will cause the substation to be inoperable and adequate local monitoring and control shall be maintained. A failure of any component should not result in an undetected loss of functions nor multiple and cascading component failures. [4]

This paper addresses process control systems and operator workstations in applications that include, but are not limited to, electric power substations. Use of these standardized reliability metrics as acceptance criteria for decisions in system design and device replacement allows direct comparison of computers from different manufacturers.

IEC 61850-3 Section 4 summarizes the design practices and reliability measures by prescribing the following quality metrics for comparison:

- Reliability measured as MTBF.
- Device availability measured as a percentage of availability.
- System availability measured as a percentage of availability.
- Device maintainability measured as MTTR.
- System maintainability measured as MTTR.

V. Hardened Computers Provide Maximum Uptime in the Office and in Harsh Environments

System and application availability should not be confused with product reliability. A relatively available system can be constructed from redundant devices with low reliability. However, this requires the complexity of redundant logic, devices, and communication, as well as the additional and constant activity of frequently replacing failed low-reliability devices, such as common low-cost personal computers for operator workstations. This has a direct and proportional negative impact on the operations and maintenance schedule and budget. Further, to maintain an available system, another device must be installed in a redundant fashion to function during the time the original device is failed and/or being replaced.

The annual computer failure rate for a specific brand of computer is simply the total number of computer failures over the last year divided by the total number of that brand of computer in service in control systems over the last year. Though this number is the most important element of choosing a replacement operator computer, many companies that specialize in office computers rather than control system computers do not know this information. Effective computer selection requires actual observed failure rates, which are provided by manufacturers that specialize in hardened control system computers. The MTBF of the computer is the inverse of the failure rate and is calculated as the total number of computers in service over the last year divided by the total number of computer failures over the last year. This measure is identified by manufacturers that track all computer shipments and returns. MTBF is also proportional in statistical terms to the working lifetime of a computer—the higher the number, the more reliable the product, and the longer the service life. The numbers are most accurate from manufacturers that do not charge to repair or replace failed computers. Choosing computers from this type of manufacturer also has the added benefit of encouraging operators to replace failed equipment with no questions asked, regardless of the source of failure. Most commercial off-the-shelf (COTS) computer manufacturers do not provide MTBF values, but the industry average is 3 years. The majority of computer failures are due to hard drives, power supplies, and fans. Hardened computers are designed for very harsh environments, so even in the relatively mild environment of the operator workstation, rugged power supplies and the lack of moving parts provide a much longer trouble-free service life. Hardened computers designed for harsh environments, such as those with temperature ranges from –45° to +70°C, have been shown to reach MTBF values well over 250 years.

For the control system industry, it is generally assumed that a replacement computer is stored on-site as part of a disaster recovery program, so the industry average MTTR is 48 hours to retrieve and install a replacement. Unavailability is the ratio of MTTR to MTBF or, in this case, 2 hours divided by the computer MTBF. Two hours is 0.000228 years, so the unavailability of the COTS computer is this value divided by 3 years, or 0.0000761. Unavailability for a hardened computer with a 250-year MTBF is 0.00000091. Availability is calculated as the unavailability subtracted from 1, so the COTS computer has an availability of 0.9999239, and the hardened computer has an availability of 0.9999991. Percent availability is calculated as the unavailability subtracted from 1 multiplied by 100 percent. The COTS computer has a percent availability of 99.99239 percent, and the hardened computer has an availability of 99.99991 percent.

These values are used to provide more actionable information during the replacement computer selection process. For example, MTBF predicts how many computers out of 250 in a large control system will fail and require repair or replacement during a given year. Annual defects per 250 installed computers are calculated as the computer population multiplied by the failure rate. When using COTS computers, operators can expect to repair or replace 83 computers out of a population of 250. When using the hardened computer, operators can expect to repair or replace one or no computers. Also, the predicted average annual downtime of a single computer is calculated as the number of minutes in the service year multiplied by the computer unavailability. For a COTS computer, operators can expect their computer to be out of service 39 minutes every year. When using the hardened computer, operators can expect their computer to be out of service for less than 1 minute, as shown in Table I.

TABLE I
COTS NONHARDENED COMPUTER VERSUS HARDENED COMPUTER

	COTS Nonhardened Computer, 2-Hour MTTR	**Hardened Computer, 2-Hour MTTR**
Unavailability	$76 \bullet 10^{-6}$	$0.91 \bullet 10^{-6}$
Availability	99.99239%	99.99991%
Average annual downtime	39 minutes	<1 minute
Computer MTBF	3 years	250 years
Annual replacements per 100 computers	83	1

VI. NEITHER SUCCESSFUL OPERATORS NOR HARDENED COMPUTERS HAVE AN OFF SWITCH

Many conclusions can be drawn from the reliability evaluation of COTS versus hardened computers for operator workstations. Although 99 percent availability sounds impressive, it actually represents a significant amount of unavailability. For example, 99 percent availability actually represents unsafe drinking water from the sink in your home 15 minutes of every day, more than 5,000 incorrect surgical operations per week, more than 200,000 incorrect prescriptions per year, and two failed aircraft landings at major airports each day. Operator workstation computers with 99 percent availability are out of service 7 hours every month.

Choosing hardened computers maximizes uptime in harsh environments, reduces failures and need for replacements, and provides more longevity than typical COTS computers. A hardened computer in the initial design reduces the likelihood of failure. Also, the choice of a hardened computer as a disaster or continuity replacement maximizes availability of the workstation and reduces maintenance cost. Perhaps most important of all, hardened computers maximize the ability of the workstation to provide the operator with information, alarms, and process control.

Reliability is the absence of failure over time. Choose a replacement computer with proven longevity and high reliability, evident as absence of failure measured as a high MTBF. Like the other control system components, computers used as operator stations in substation and industrial control systems have no off switch. Once these devices are plugged in and turned on, they are expected to run trouble-free and uninterrupted for years on end.

VII. DO NOT CONFUSE FAULT-RESISTANT WITH FAULT-TOLERANT DESIGNS

Operator console hardware that is designed to be fault tolerant requires that a failed computer be replaced with a new computer while the system remains operational. This is often referred to as hot-swapping the failed computer. Operator workstations with a single in-service backup computer are known as single-point tolerant. In these systems, the MTBF needs to be high enough for the operator to have time to replace the failed computer (MTBF greater than MTTR) before the backup also fails. As long as operators are available to constantly replace failed computers and the computer MTBF is greater than the MTTR, the operator console is fault tolerant. Of course, this elevates the likelihood of losing information and alarms during the swapping process and drastically increases the cost and complexity of maintenance.

Carefully designed fault-tolerant applications based on high-MTBF computers can actually be categorized as fail-safe or nonstop systems. The advantages of dual computers are numerous; the disadvantages are numerous as well. The most obvious disadvantages are the complexity of the failover scheme and the cost of running parallel computers.

Fault-resistant designs are much simpler to implement. They simply require that the computer exhibit high availability, measured as a high MTBF, and that, should the computer fail, its impact on the system be quarantined. In this case, the system continues to run the control system processes during the time it takes the operator to replace the computer (MTTR).

VIII. QUARANTINE FAULT-RESISTANT COMPUTER PROCESSES FROM NONSTOP SYSTEMS

Process control or protection, control, and monitoring systems are designed to be fault resistant; however, the operator workstations within these systems are typically designed as fault resistant rather than fault tolerant. If the operator computer fails, the system continues to run in a diminished capacity. The system may not fully collect information, logs, and alarms without the operator computer. However, correctly designed control systems have decentralized locations for operators to perform commanded control during the short time it takes to replace the computer. In this case, the processes to collect, store, and report logs, statuses, and alarms in the operator computer are quarantined from the rest of the system. In this way, these processes are isolated, suspended, and then restored without significant impact on the rest of the control system.

The advantages of fault-resistant operator workstation designs are numerous. They are much simpler than dual-computer designs, they are much less expensive, and high-MTBF computers mean long periods of trouble-free operation. Also, it is quite easy to quarantine the interface of a relatively slow-acting operator and let the fault-tolerant design run every few milliseconds to protect, control, and monitor systems. In fact, the only disadvantage is not obvious but directly impacts the operator. Even if the operator is aware of the fault, having a fault-tolerant system can inadvertently confuse the operator about the importance of repairing the fault. If the faults are not corrected or if the computer is replaced with a low-MTBF substitute, this eventually leads to system failure when the fault-tolerant component fails completely or important logs and process alarms go undetected without visibility on the operator computer.

A catastrophic example of the impact of a failed operator workstation in the news recently is the tragedy on the Deepwater Horizon oil platform in the Gulf of Mexico. An article reporting on a congressional hearing about the failed alarming and control systems illustrates the direct result of failed computers and processes [5]. At a federal hearing in April 2010, the chief electronics technician on the oil platform, Michael Williams, said "that the rig's safety alarm had been habitually switched to a bypass mode to avoid waking up the crew with middle-of-the-night warnings" [5]. He described how the partially fault-tolerant computer system continued to record high gas levels or the presence of a fire. However, this created a false sense of security because the system "would not trigger warning sirens" [5]. Williams described how he had been asked to check the computer system that monitored and controlled the drilling five weeks before the April 2010 explosion. He found that the machine had been locked up for months and "you'd have no data coming through" [5]. With this failed computer workstation, the drillers had no access to "crucial data about what was going on in the well" [5]. Eleven people died in the disaster, which was the largest oil spill in United States history.

It is the responsibility of the operator to make sure that the computers are working correctly by doing the following:

- Reporting failed equipment and processes.
- Reporting when a computer locks up and shows a blue screen.
- Finding someone to tune alarms to be relevant and informative and change nonalarm inputs to status.
- Finding someone to repair or replace failed equipment.

IX. Physically Protect Operations Computers From Attack or Mischief

Though it is possible to interpret the physical security of the control room or substation control house as the physical security of a computer workstation, it is not appropriate. The operator must maintain perimeter barriers that include both physical and psychological deterrents to unauthorized entry, delaying or preventing forced entry. The most secure method is in fact the six-wall approach, recommended to secure critical cyberassets.

The U.S. Federal Energy Regulatory Commission (FERC) Order 706, issued January 18, 2008, and recently reaffirmed on January 21, 2010, approves eight standards to maintain mandatory reliability requirements for Critical Infrastructure Protection (CIP) [6]. The CIP reliability standards require certain users, owners, and operators of the bulk electric system to comply with specific requirements to safeguard critical cyberassets. These standards provide useful guidance for performing due diligence planning and execution of security practices, regardless of the system association with the bulk electric system.

The purpose of North American Electric Reliability Corporation (NERC) CIP-006-1 is to provide guidance on methods to physically protect critical cyberassets and those devices used in the access control and monitoring of physical security perimeters and electronic security perimeters. CIP-006-01 describes that a physical security plan must include an electronic security perimeter and a physical security perimeter. As defined in the NERC glossary, an electronic security perimeter is "the logical border surrounding a network to which Critical Cyber Assets are connected and for which access is controlled" [7], and a physical security perimeter is "the physical, completely enclosed ('six-wall') border surrounding computer rooms, telecommunications rooms, operations centers, and other locations in which Critical Cyber Assets are housed and for which access is controlled" [7].

For a computer, the six-wall border is a physical, completely enclosed border, such as a room, cage, safe, or metal cabinet. Raised floors and drop ceilings do not constitute part of a border because they create potentially uncontrolled access points. Access to the keyboard, mouse, communications ports, and interface cables creates unsafe access points when the room is not secure. Therefore, a six-wall solution is to mount the computer into a locking metal cabinet with cabling passing through a grommeted hole. This prevents unwanted access to any unused access ports, such as USB, Ethernet, or serial ports. It also prevents someone from unplugging a legitimate cable connection and hijacking it for unauthorized access. The keyboard and mouse should be in a retractable and lockable tray, with the key provided only to operators.

A public example of the impact of a compromised USB connection in the news recently is the attack that repurposed a keyboard connection into unauthorized access of a computer. The Black Hat conference series markets itself as a forum for individuals to provide cybersecurity information in a vendor-neutral environment. A news article about a Black Hat conference describes a public demonstration of unauthorized access of a computer through a repurposed USB keyboard connection [8]. The hackers unplugged a keyboard, connected a smartphone in its place, and posed as the keyboard to gain access to the computer. This hack, which can be staged locally and performed remotely, "takes advantage of USB's inability to authenticate connected devices coupled with the operating systems' inability to filter USB packets" [8].

This further illustrates the need to physically secure computer access ports because even those in use for computer peripherals can be hijacked for unauthorized access.

A public example of the impact of compromised removable media peripherals in the news recently is the attack on the Pentagon computer network. A news article about the theft of information via an unsecured DVD peripheral describes how the information was made available to the WikiLeaks organization [9]. Because of the breach that released thousands of classified documents from the Pentagon's secret network, the U.S. military is telling troops that they risk court martial if they do not stop using CDs, DVDs, thumb drives, and every other form of removable media. Pfc. Bradley Manning admitted to downloading "hundreds of thousands of files … to a CD marked 'Lady Gaga'" [9] and then gave the files on CD to the Wikileaks organization. Further illustration of the need for diligence at each computer is the fact that this occurred shortly after a ban on using removable media was lifted by the Pentagon. The ban had been in effect for two years after removable media had been responsible for spreading a worm to hundreds of thousands of computers. A worm cleanup effort was completed, and the ban was lifted in February 2010. Soon after, Manning began providing information to Wikileaks.

The Pentagon example shows how removable media were used to both introduce malware and steal classified information via nonsecure computer access in secure control rooms. The previously mentioned Night Dragon incident demonstrates a specific attack on data stored on operator computers in process control systems. The hackers used two external methods to compromise computers—through public websites or through infected emails sent to company executives. The hackers focused on documents detailing oil and gas field exploration and bidding contracts, as well as proprietary industrial processes. This information is tremendously sensitive and worth a large amount of money to competitors.

Both the Pentagon and Night Dragon examples point out the need for operators to restrict physical access to their computers, even inside a secure control room.

X. Do Not Become a Black Hat by Bypassing Whitelisting

Many control systems are separate, compartmentalized networks. The best practice is to not connect these systems to the Internet and to tightly control access to internal business networks [10]. In these cases, companies spend a lot of money and effort securing the networks. With these best practices (such as whitelisting) in place, laptops and USB keys that connect to these networks are the main way that unwanted and malicious software, or malware, is introduced into this environment. Whitelisting essentially documents all appropriate applications that are allowed to run on a computer and helps to block others from running if they are inappropriately added or modified in the computer.

USB and other removable media are real attack vectors that can often be overlooked—especially in networks protected by other means of isolation. It is no mistake that the most sophisticated worm ever written spread through USB. The writers of Stuxnet knew that the control systems would be isolated from the Internet and that they could not attack nuclear networks directly through public communications connections. Instead, they must have realized that USB would be an effective way to evade protection. They are not the only ones. In fact, 25 percent of new malware in 2010 had a component to spread through USB [11]. Properly configuring Microsoft® Windows® Autorun behavior; utilizing technologies such as whitelisting, regularly updated antivirus, and device control; and limiting administrative rights on control systems will help reduce the risk of USB-spread malware—but do not just assume those controls are enough [12] [13]. There are additional things you can do to help. Visual reminders, such as the poster shown in Fig. 1, encourage people to be cautious when using these devices, particularly in control environments.

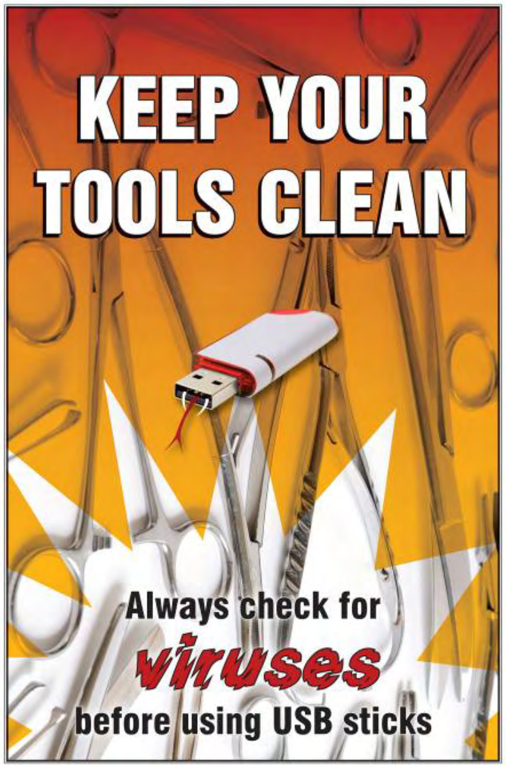

Fig. 1. Convenient USB storage devices can spread malware.

Just like control networks are designed with limited access based on need, USB use should be limited and based solely on the needs of the system. In fact, if USB storage is not needed on Windows computers, USB interfaces can and should be disabled natively in Windows [14]. It is best to choose computers that allow the operator to disable each peripheral and connection when not in use. Personal, casual USB storage devices have no place on a critical network where many other methods of protection have been put in place. Additionally, those casual devices are often at a higher risk of infection. Many people do not realize that Windows loads the storage component of an MP3 player or mobile phone plugged into the USB port, even if the only intent is charging a low battery. The same is true of plugging in through USB hubs that expand the number of USB ports on a computer, such as those on some monitors and keyboards.

If there is a specific, authorized need to connect via USB to collect logs or perform updates, consider the analogy of performing surgery on your system and always use a sanitized device. On the typical enterprise computer, Autorun is disabled. For performance reasons, most antivirus solutions do not perform a full scan of USB devices. It is possible to plug an infected USB key into a Windows computer with fully updated antivirus, copy additional files onto the USB key, and remove it—without Windows ever detecting the malware! For safety, it is recommended that USB storage devices be sanitized by either formatting or at least fully scanning the USB storage device for malware before connecting it to any protected control system (to do approved work). Do this before or after connecting to any unknown system, including between connections to multiple systems that do not have whitelisting or current antivirus protection. In fact, best practice is to fully sanitize each USB device after each use.

XI. Do You Have Ways to Remember Security Tips?

A simple way to remind operators of the dangers of using removable media is to insert warning plugs in all empty USB ports and warning disks in CD and DVD drives. These devices carry reminder warnings and require the operator to physically and consciously remove them to gain access to the computer. This helps remind operators that it is their responsibility to secure these ports and to resecure them after use. Another simple way to raise and maintain awareness is to post warnings and cybermessages in the workplace, similar to other workplace-related reminders. These are available in the form of sensible cybersecurity best practice posters [15], and the United States Computer Emergency Readiness Team (US-CERT) has a section dedicated to control system best practices for review [16]. Public awareness reminds everyone that it is each person’s responsibility to maintain and secure operator computers.

XII. Observe and Analyze Alarms and Alerts to Detect Problems Quickly

Alarms and log files provide invaluable information but can be overburdening if not set up correctly. Knowing the location of your alerts and log files is important. In the case of the Deepwater Horizon explosion, fire alarms were disabled months before the explosion occurred [17]. This type of tampering is more common than you might think. When a system continues to report false positive alarms, people tend to either disregard the events or shut them off. The best approach is to find root cause for the false positive events and take corrective action—it might just require some simple tuning of thresholds. Another problem is the volume of events that can occur on a system. It is usually not feasible for one person to review all the logs on a system, so certain actions must be taken prior to major events occurring. One approach is to review your logs periodically and know what is normal and what is not normal for your system. In the case of normal events, they can usually be filtered out, and you can focus on the high-severity alerts. Another approach is to use a central logging mechanism that allows you to make correlations on these events. If a failed login attempt occurred, it may not be that important. However, if many failed login attempts occurred directly followed by a successful login event, it may be more significant because it may be a brute force password-guessing attack against your system. In any case, alarms must be tuned and cannot be ignored. Visual reminders, such as posters, help people remember how important it is to monitor logs, alarms, and reports.

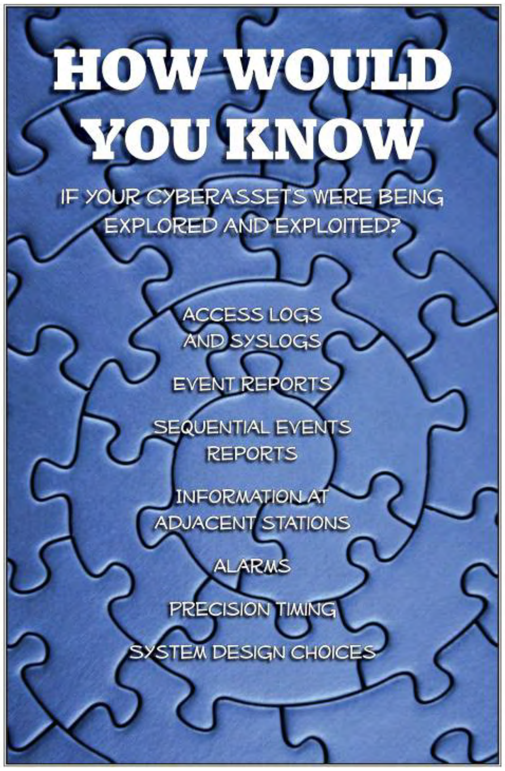
HOW WOULD YOU KNOW
IF YOUR CYBERASSETS WERE BEING EXPLORED AND EXPLOITED?
ACCESS LOGS AND SYSLOGS
EVENT REPORTS
SEQUENTIAL EVENTS REPORTS
INFORMATION AT ADJACENT STATIONS
ALARMS
PRECISION TIMING
SYSTEM DESIGN CHOICES

Fig. 2. Looking at the warning signs—including logs, alarms, and reports—is critical in detecting and responding to problems promptly.

XIII. Do Not Hand Over the Keys to Your Systems

There are many reasonable measures that can be utilized to prevent unauthorized access to computing systems. The first and easiest measure is to log out or lock your system when you leave it unattended. You might have a very strong password, but if you leave your system unlocked while unattended, you are setting yourself up for the risk of unauthorized actions with your credentials. Not only is this a security risk, but it could also pose a compliance violation.

Second, create your account with the least privileges possible to get your job done. Granting full administrative access to your account when you do not need it is both unsafe and a way that systems are compromised by malware.

Another avenue of protection against unauthorized access is securing your credentials properly. In many systems, passwords are used to allow and deny access and thus must be sufficiently strong to prevent attacks. Do not share your password with others or leave your password on a piece of paper. Make sure your password is long enough so it cannot easily be guessed. Do not use passwords based solely on dictionary words. Apply complexity rules to your password, using uppercase and lowercase letters, numbers, and special characters, as shown in Fig. 3. Enable account lockout and logging for consecutive failed password attempts, which prevents and helps locate where unauthorized access originated. Never reuse passwords or use the same password for different systems; otherwise, it only takes one compromised system for all other systems to be vulnerable. HBGary, a security company, recently had multiple systems exposed because their administrators used the same passwords for multiple purposes [18]. The hackers gathered website passwords and used those passwords to connect to email, social networks, and other resources. Use strong, unique passwords for different systems and different purposes.

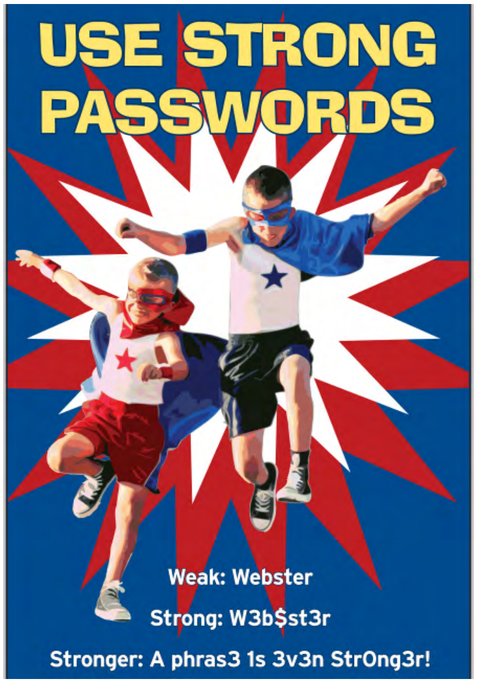

Fig. 3. As attackers become more sophisticated and have more computing power to leverage, strong passwords become more critical. Often, passphrases are the strongest combination of easy to remember and hard to crack.

XIV. Do You Know What Normal Looks Like in Your Environment?

One of the keys to maintaining a secure system is knowing what normal looks like. Some types of data to include in a baseline might include:

- Standard connected equipment
- Communications ports and protocols
- Communications paths
- Typical load, alarms, and logs
- Security settings
- Files

Attackers always leave a trace. Stuxnet added mysterious files to USB keys, Night Dragon disabled antivirus, common Structured Query Language (SQL) injection attacks generate excess logs, and bots initiate unusual network traffic. In most cases, it is a combination of unusual events. Are you looking for those anomalies? Read [19] for more information about identifying anomalies and tracing them back to an attacker.

XV. Am I Ready to Respond to an Incident?

Practicing good security, including the tips in this paper, will help protect your systems and reduce the risk of attackers infiltrating them. Even with the tightest security in place, it is important to plan and prepare for recovery. When companies are not prepared for an incident, it can lead to slower recovery times, loss of criminal evidence, accidental disclosure of sensitive information, and public relations problems. Before an incident happens, ask yourself:

- Do I know when to declare or escalate a potential incident?
- How would I recognize a security incident? [19]
- Do I know what is expected of me?
- Who do I alert and when?
- Do I know the recovery procedures?
- Are the procedures different if foul play is suspected?

As they say, an ounce of prevention is worth a pound of cure. Be thinking of ways you can better protect your systems from an incident, and apply them now before an incident happens. In addition, make sure you have measures in place to help detect anomalies in the system that might be an indication of an intrusion. Lastly, if an incident occurs, identify root cause and learn from it. There are always opportunities after an incident to look back and see what you could have done differently, how your systems could have been configured differently, and which processes need improvement.

XVI. Summary

As control systems become more of a target and more interconnected with business systems, security will continue to become increasingly critical. Ignoring or disabling alerts or logs, carelessly using removable devices, using unauthorized and nonsecure network connections, leaving systems unlocked, not addressing other physical security concerns, or failing to practice basic security (using weak passwords) puts your control systems at risk.

It is your responsibility to know the baseline performance, behavior, and attributes of your operator workstation. It is critical to understand what your systems should look like and what normal is so you can detect differences and identify problems.

It is your responsibility to negotiate with authorized personnel to tune alarms and reporting to be accurate. Prevent the system from creating nuisance notifications that lead to operators assuming all alarms are trivial.

It is your responsibility to work with appropriate personnel to make changes to whitelisting and/or add new applications. It is essential that ongoing use of the computers meets the security requirements of the system and corporation.

It is your responsibility to choose high-MTBF computers as initial and replacement operator workstations. Minimized downtime improves the financial return on investment but, more importantly, provides better operator safety, peace of mind, and success.

It is also clear that no single method mentioned here is completely successful. Each failure creates or increases your system vulnerability. Multiple failures, like decreased awareness due to poor alarming in conjunction with a failed operator computer, increase system vulnerability exponentially.

There have been catastrophic events in the news because technicians disabled or ignored alarms. Intruders are counting on you to take shortcuts in securing your systems and use dirty USB devices to spread their malware, while betting that you will not notice. It is up to you to prove them wrong.

System and security planners are responsible for the security and reliability design of the control system. However, it is the sole responsibility of the operator at the keyboard to take these steps to ensure the safe, reliable, and economical performance of the control system.

XVII. References

[1] "Chinese Hackers Hit Five Major Energy Firms," Fox News, February 2011. Available: http://www.foxnews.com/scitech/2011/02/10/chinese-hackers-hit-major-energy-firms.

[2] E. O. Schweitzer, III, "Ten Tips for Improving the Security of Your Assets," November 2009. Available: http://www.selinc.com.

[3] IEC 61850 Standard. Available: http://www.iec.ch.

[4] IEC 60870 Standard. Available: http://www.iec.ch.

[5] G. Keizer, "Tech Worker Testifies of 'Blue Screen of Death' on Oil Rig's Computer," *Computerworld*, July 2010. Available: http://www.computerworld.com/s/article/9179595/Tech_worker_testifies_of_blue_screen_of_death_on_oil_rig_s_computer.

[6] FERC 18 CFR Part 40, Order No. 706, *Mandatory Reliability Standards for Critical Infrastructure Protection*, January 2008. Available: http://www.ferc.gov.

[7] North American Electric Reliability Corporation, "Glossary of Terms Used in NERC Reliability Standards," *Reliability Standards for the Bulk Electric Systems of North America,* April 2010. Available: http://www.nerc.com/files/Reliability_Standards_Complete_Set.pdf.

[8] M. Gorman, "Hackers Disguise Phone as Keyboard, Use It to Attack PCs Via USB," *Engadget*, January 2011. Available: http://www.engadget.com/2011/01/23/hackers-disguise-phone-as-keyboard-use-it-to-attack-pcs-via-usb.

[9] N. Shachtman, "Military Bans Disks, Threatens Courts-Martial to Stop New Leaks," *Wired*, December 2010. Available: http://www.wired.com/dangerroom/2010/12/military-bans-disks-threatens-courts-martials-to-stop-new-leaks.

[10] T. Nash, "An Undirected Attack Against Critical Infrastructure: A Case Study for Improving Your Control System Security," US-CERT Control Systems Security Center Case Study Series: Volume 1.2, September 2005. Available: http://us-cert.gov/control_systems/practices/documents/CaseStudy-002.pdf.

[11] "25 Percent of New Worms Designed to Spread Through USB Devices," *SecurityWeek News*, August 2010. Available: http://www.securityweek.com/25-percent-new-worms-designed-spread-through-usb-devices.

[12] Microsoft Support, "How to Disable the Autorun Functionality in Windows," September 2010. Available: http://support.microsoft.com/kb/967715.

[13] D. Anderson, "Increase Security Posture With Application Whitelisting," proceedings of the 13th Annual Western Power Delivery Automation Conference, Spokane, WA, March 2011.

[14] Microsoft Support, "How Can I Prevent Users From Connecting to a USB Storage Device?" September 2009. Available: http://support.microsoft.com/kb/823732.

[15] Sensible Cybersecurity Best Practices Posters. Available: http://www.selinc.com/cybersecurity/posters.

[16] United States Computer Emergency Readiness Team, Control Systems Security Program, Recommended Practices. Available: http://us-cert.gov/control_systems/practices/Recommended_Practices.html.

[17] "Deepwater Horizon Alarm System Was Partly Disabled Prior to Explosion, Technician Tells Congress," *The Huffington Post*, July 2010. Available: http://www.huffingtonpost.com/2010/07/23/deepwater-horizon-alarm-s_n_657143.html.

[18] B. Zdrnja, "HBGary Hack: Lessons Learned," Internet Storm Center Diary, February 2011. Available: http://isc.sans.edu/diary/HBGary+hack+lessons+learned/10438.

[19] E. O. Schweitzer, III, D. Whitehead, A. Risley, and R. Smith, "How Would We Know?" proceedings of the 37th Annual Western Protective Relay Conference, Spokane, WA, October 2010.

XVIII. Biographies

David Dolezilek received his BSEE from Montana State University and is the technology director of Schweitzer Engineering Laboratories, Inc. He has experience in electric power protection, integration, automation, communication, control, SCADA, and EMS. He has authored numerous technical papers and continues to research innovative technology affecting the industry. David is a patented inventor and participates in numerous working groups and technical committees. He is a member of the IEEE, the IEEE Reliability Society, CIGRE working groups, and two International Electrotechnical Commission (IEC) technical committees tasked with global standardization and security of communications networks and systems in substations.

Bryan MacDonald is the director of security at Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington. He has 15 years of experience working in information technology with an increasing focus on security. He has worked in help desk, system administration, infrastructure management, and data security prior to taking on his current role. He enjoys helping people understand security risks and providing corresponding sensible security recommendations and support. He is a Certified Information Systems Security Professional (CISSP) and has served as the outreach director for the Palouse chapter of the Information Systems Security Association (ISSA). Bryan studied computer science and mechanical engineering at Washington State University. He is also an avid cyclist and gadget lover.

Jason Kraft received his BS in Computer Science from Washington State University. He is now a data security analyst for Schweitzer Engineering Laboratories, Inc. (SEL) in Pullman, Washington. Prior to joining SEL in 2005, he worked as an independent consultant, working on projects such as designing and implementing network management tools, network auditing, and system administration. Previous to that, he worked for Amazon.com as a network tools and analysis engineer, where he implemented software to assess the network infrastructure. He holds the certification of Certified Information Systems Security Professional (CISSP). He is a current member of the Spokane and Palouse Information Systems Security Association (ISSA) chapters.

Peter Dolezilek is a high school student in Pullman, Washington, working part time at Schweitzer Engineering Laboratories, Inc. He works in the security department creating and distributing materials to promote security awareness throughout the company.

20110314 • TP6494-01