Cybersecurity Best Practices for Creating Resilient Control Systems

Jess Smith, Joshua Pereyda, and Dennis Gammel
Schweitzer Engineering Laboratories, Inc.

This paper was presented at Resilience Week 2016 and can be accessed at: http://dx.doi.org/10.1109/RWEEK.2016.7573308.

For the complete history of this paper, refer to the next page.

Published in
Sensible Cybersecurity for Power Systems: A Collection of Technical Papers Representing Modern Solutions, 2018

Originally presented at
Resilience Week 2016, August 2016

Cybersecurity Best Practices for Creating Resilient Control Systems

Jess Smith, Joshua Pereyda, and Dennis Gammel
Schweitzer Engineering Laboratories, Inc.
Pullman, WA USA
jess_smith@selinc.com

***Abstract*—Control systems manage the automated systems that run our world. The security of these systems is critical, but most modern best practice guidelines focus on enterprise security, similar to that of traditional information technology and security. This paper provides eight basic guidelines to increase cybersecurity in control systems.**

***Keywords*—control system; SCADA; EMS; ICS; cybersecurity; best practices.**

I. Security in Control Systems Is Critical

Control systems surround us in our homes and workplaces; move us in our cars, trains, and airplanes; and enrich our world through the production and transportation of electric power, food products, fossil fuels, and critical chemicals. In the modern age, humans are too slow and don't have the endless patience needed to constantly monitor the many inputs and outputs required to keep all of our manufacturing floors, substations, or building heating/cooling systems running. We may call it a supervisory control and data acquisition (SCADA) system, an energy management system (EMS), a building automation system (BAS), a distributed control system (DCS), or one of many other names, but these are all subtypes of control systems, and they all are vital to modern life. The way in which these systems interact with the real world make them valuable targets for hackers, in dire need of security. [1]

The best way to ensure your control system network is cyber secure is simply to take it off the Internet. If there are no Internet connections, no cyber ways for viruses or hackers to get in, the network is much less likely to be compromised. It is not impossible for something bad to get in; the prevalence of removable media like USB sticks still provide an intrusion vector, but the risk is reduced by orders of magnitude. We strongly discourage connecting any control systems to the Internet. However, either by accidental connection of control systems to the internet or through misguided efforts to improve usability and efficiency, some control systems are connected to the Internet. It may be reasoned that it is simply much easier to remotely log in to the substation or pump house and fix the problem, rather than driving out to do the same [2]. When evaluating our control system security let's consider Internet connection a possibility and our worst case scenario. Worst case is what we need to be planning for, and if we plan for worst case we create the best case scenario.

This paper provides and discusses the following eight best practices for designing, building, and securing control systems:

1. Know, limit, and monitor access to the control system.
2. Implement the appropriate security for each level of the control system.
3. Continuously monitor the control system at all levels.
4. Have a contingency plan.
5. Patch, update, and maintain.
6. Don't forget physical security.
7. Learn from events.
8. Be aware of your public information.

II. Best Practices

A. Know, Limit, and Monitor Access to the Control System

Your control system is your most valuable resource. Without it, power is not created or distributed, the manufacturing floor shuts down, or the airplanes won't get off the ground. People should not be able to walk in off the street and access the system with a default password. The users who do need to access it should be vetted, limited to what they actually need to use, and then held accountable for their actions. [3]

When considering risk, one of the first thoughts any control system defender must have is, where are all the places an intruder can get into the network? Internet connections and physical access are two of the most obvious, but USB drives, microwave and other wireless communications, and power supplies also need to be considered and protected. In a modern control system network, it is entirely possible to have access points down at the intelligent electronic devices (IEDs), via USB, so be sure to consider the whole network, not just the desktop PCs.

Default passwords on devices are found frighteningly often on systems in use. These default passwords were set at the factory, all the same, to make it easier for end users to

customize the device. The intention is that the end user will change the default password to something unique and strong, but many end users, for ease of use, leave the password as the default. Default passwords are also widely available on the Internet, both from vendor's or manufacturer's websites and online manuals as well as from compiled lists of default passwords on third-party websites. It is safe to assume that any default password can be found somewhere on the Internet.

Similarly, avoid common words or number sets. There is a list that comes out each year, covering the top 25 most popular passwords. Repeatedly, year after year, passwords like "123456," "password," "qwerty," "baseball," "football," "Yankees," "Steelers," and "Lakers" are found on the top of that list. Common numbers or words are not cryptographically secure. Passphrases, substitution, and slang are all better ideas for creating stronger passwords. For example, you could use something like "S,tff:2bgwnmhgb!" (Space, the final frontier: To boldly go where no man has gone before!). This is not a dictionary word, it contains letters, numbers, and symbols, and it is nice and long yet memorable. A reminder, though—never use a password or passphrase that you have seen online or in a paper (including this one!). Hackers read these papers, too.

There are three "As" that are important in ensuring that access to the control system is limited—authentication, authorization, and accountability. We first need to authenticate the user, to ensure that they are who they claim to be. This can be done through verifying one or more of three things about the user: what they know (a password or PIN), what they are (biometrics), or what they have (a card or token).

After we have ensured that Technician Joe Snuffy is actually Joe Snuffy (authentication), we need to limit the areas Joe can access. As a technician, Joe should be able to view and change some device settings, but should not have full engineering access to the entire network. This is authorization. Finally, we also need accountability, so that if Joe goes in and changes settings, there is be a record of that. In case of a problem, this allows us to both educate users to prevent future mistakes as well as discover the entry point for a malicious intruder.

B. Implement Appropriate Security at Each Level

Not all parts of the ICS are created equal, and not all the devices in the control system should be secured identically. A defense-in-depth method has been proposed by the United States Department of Homeland Security [4], Schweitzer and Frincke [5], and many other researchers, which we have modified and expanded for related work. A precis of this can be seen in Fig. 1.

A control system network has five levels, each with its own considerations for security and monitoring. We start the levels at the bottom, Level 0. Level 0 is comprised of basic sensors without any decision-making capabilities embedded in them. These sensors need to be physically protected, should be redundant to prevent failure-based errors, and need physically protected communications lines. Moving above Level 0, we begin to see the digital controllers coming into play. Levels 1 and 2 include digital controllers, with the line between the two being that Level 1 contains serial-based communications and Level 2 contains Ethernet-based communications. It is possible to have devices that straddle the line between these two layers. This is a critical distinction, however, because how we monitor and protect serial communications (usually point to point) is very different from how we monitor and protect Ethernet communications (routable). [6]

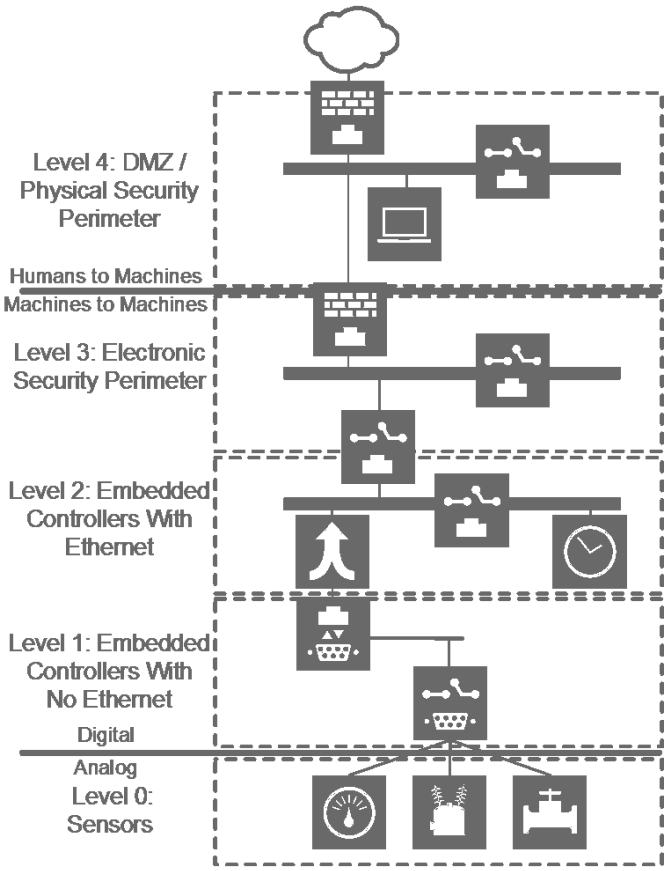

Fig. 1. Control System Network Levels

Level 3 is the top of the machine-to-machine communications layers and includes the electronic security perimeter. There is a division here to highlight the difference between the machine-to-machine and the human-to-machine communications. While machines are repeatable, consistent, and largely deterministic, humans are nondeterministic and cannot be predicted in the same way that a machine can. Because of this, techniques like whitelisting are possible at Levels 3 and lower, while blacklisting is a better option for Level 4. Level 3 is also critical because it marks the boundary between critical and optional functionality. Our control systems should be able to continue functioning without direct human intervention for periods of time, and Level 3 needs to protect those crucial functions from the greater risks found at Level 4.

At Level 4, with the humans and regular PCs with normal operating systems, we will have to deal with a huge range of challenges, including but certainly not limited to malware, phishing, user errors, and code bugs. A DMZ, or demilitarized zone, between the humans and the critical functions at Levels 3 and below is a valuable barrier.

1) Evaluate Risk for Each Level and Component, and Implement the Appropriate Security

Each of the individual network levels shown in Fig. 1 has its own strengths and weaknesses and needs to be evaluated separately and as part of the whole. A cyber attacker is much more likely to be able to penetrate and gain access to Level 4 than Level 1, and the methods an attacker would use to reach those levels are very different. Alternately, it may be much easier for a physically present attacker to gain access to the Level 0 sensors placed outside of the physical security perimeter. Once we have evaluated the risks and found the weaknesses for each level, we can begin to protect against and mitigate those risks.

2) Segment the Network

Networks should be segmented. In the same way that a human user does not always need to access the whole network, the elements of the network do not all need to be able to access each other. By segmenting the network, we are limiting both human and machine traffic to only necessary areas. This prevents infections or breaches from spreading as easily. Segmentation from the larger network can be easily provided at Level 3, but there should also be further segmentation within Levels 1 and 2. In the same way that we segment our regional power grids to create survivable islands, we should also segment our control system networks.

C. Continuously Monitor the Control System at All Levels

Security is ongoing. Continuous monitoring of the network is critical to catch intruders or infections, and monitoring should be performed at all levels of the network. The different levels will show different data, and sometimes we can see problems by combining the data from multiple levels which are not obvious in a single level. It is also necessary to introduce automatic monitoring and automatic log reviews as the amount of data to be processed and considered rapidly grows to be larger than any human can manually parse. [7]

1) Baselines

Baselining is a critical part of monitoring; how do we know what is wrong, unless we have a baseline to tell us what is right? Baselines should include as many data points from as many different devices (network devices, sensors, and controllers all) as possible, and should stretch over enough time to detect regular patterns. For example, in a normal office environment, users will be logging in Monday to Friday, between 7 a.m. and 6 p.m. A baseline that only looks at one weekday will then throw a false positive error on Saturday, when no one logs in.

2) Alarms

Humans don't do well at continuous monitoring by themselves. They get distracted, bored, or tired. This is a huge part of why we have control systems in the first place; the machines are always paying attention. However, there are times that the machines hit something they cannot deal with or a situation that is outside of the allowable parameters. Alarms are the most direct, immediate way that a control system has to interact with the human operators. Alarms should be pre-programmed to trigger whenever the normal operation of the control system ventures outside of the pre-programmed and acceptable bounds. Traditionally, this was a flashing light or annoying sound, but modern control system can also use emails and text messaging to alert the operator that something is not quite what it should be.

3) Logs

Logs are a valuable tool for figuring out what happened after something goes wrong. The alarm has gone off; the boiler is overheating. After we deal with the immediate problem, it is critical to figure out how and why the boiler got to that state and prevent it from happening again. The same applies for a cyber attack. Properly configured logs will help us to discover how far into the network the attacker got, what they modified, and how they got in in the first place. Logs can be set to record firmware updates, network traffic, access attempts/successes/failures, usage statistics for hardware like processors and memory, and many other things custom to your system.

D. Have a Contingency Plan

Something will go wrong in your control system—it is not a matter of *if* but rather *when*. We can do our best to prevent the intrusion or infection, and to limit the reach and damage when they get in, but a good control system defender will always have a plan for the worst case scenario. [8]

1) Have Several Plans, One for Each Contingency

There are probably multiple worst case scenarios, and you should have plans for as many of them as you can think up. Consider both malicious and non-malicious scenarios; non-malicious scenarios can include both human-caused problems (someone hit the wrong button, downloaded a virus, etc.) as well as other acts of God.

2) Have the Plans Printed Out and Near the Control System or Control Center Where They Will Be Used

A plan that is on a computer when the power goes out, or a plan physically printed out but hundreds of miles from the problem, are not useful plans. The plan should be near wherever the fix will be implemented and printed out so that, in case of electronic failure, it can still be read.

3) Update Plans Regularly

Your control system changes, and your contingency plans should change along with them. It may not be feasible to

update the plan when every little change is made, but at least yearly the plans should be reviewed and updated to match with the control system.

4) Practice the Plans

Treat your control system contingency plan like a fire escape plan and have drills regularly to ensure everyone knows where they need to go and what they need to be doing. New employees or contractors need to be made aware of where the plans are, in case they are the only one available when a problem occurs, and what their normal duties with respect to those plans are. Experienced employees will need periodic refreshers, to maintain their awareness.

E. Patch/Update/Maintain

New threats emerge every day, and a good control system defender needs to ensure that their control system is kept up to date with defenses against those threats. Part of a responsible control system defender's job includes monitoring the news, blogs, mailing lists, and other sources to keep on top of those new threats. It is critical that we create processes to ensure that we are appropriately keeping the control system up to date.

1) Test First

Control systems are not like enterprise networks; they have a wider variety of very different components, and lives may be at stake if the network fails. Test any update out on a non-connected machine before placing it on a live system. If your budget allows for it, a simple secondary system used as a sandbox for testing updates as well as security assessments can be a valuable tool.

2) Have a Schedule for Updates

Don't let updates become a "when we get around to them" thing. An un-implemented patch or update may leave a known weakness unprotected and your control system vulnerable. Having a plan, with a schedule that you hold to, will ensure that updates are performed in a timely fashion.

3) Have a Plan for Critical Updates That Come Out of Cycle

Thankfully rarely, sometimes an exploit becomes public that is a level above the norm. It may be either very widespread, such that all systems have it, or of enormous risk, allowing an intruder access to critical systems. When this happens, a plan needs to be in place for getting a patch or update pushed out to the system as soon as possible. This does not reduce the need for testing, but it may increase the speed and priority of the testing for that patch.

4) Hardware Has a Lifecycle

We update and patch software regularly, and need to do the same for our computing hardware. The processors, hard drives, and other components in our devices will eventually fail, and a plan must be in place to detect failure. If possible, a plan to rotate or update computing hardware prior to failure should be in place. By fixing or replacing the devices before the problem hits, we can reduce downtime and prevent damage to the system.

F. Don't Forget Physical Security

Much of cybersecurity is focused on electronic penetration, coming from over the network or the internet. In a control system, however, we also must have care for our physical security. Often, the front port of a device is left with default passwords, and if an intruder can cut through the fence surrounding the control system, they can access the devices easily.

Communications paths should be physically protected as well. If the devices are behind a wall, but their trusted communication path is out in the open (either via a physical line, like Ethernet, or via an Internet tunnel), the devices are vulnerable to man-in-the-middle or passive listening attacks. Encryption can be a powerful tool to protect those communications that pass outside of our control, but care must be taken with the methods and implementation of that encryption.

G. Learn From Events

When the system goes down, for whatever reason, we often focus solely on getting the system back up and running. After the adrenaline runs out and the control system is back up, we still need to learn from what just happened. How did the intruder or infection get in? How did it propagate and move within the network? What data was exfiltrated, what settings changed, what user accounts compromised?

1) Learn From Both Your Events as Well as Other Organizations

A wise man sees another man burn his hand on the stove and chooses not to touch it himself. We should attempt to do the same thing with our control system and learn from others' mistakes and problems, to improve our networks without the actual intrusion or infection. In section E, we discussed the necessity of a good defender maintaining situational awareness, and this knowledge can be used to look forwards to prevent problems.

If your network does have an event, after you get the system back up and running, take that opportunity to look honestly and openly at why the situation occurred and learn from it. Take blame out of the picture and focus on figuring out what when wrong and how to fix it for the future.

2) Test

We should be testing our own networks and people, looking for weaknesses and vulnerabilities so that they can be fixed before the malicious users find them. Common tools such as Wireshark and Nmap can be used to explore what is being communicated over your network and what devices are available. Likewise, there are free vulnerability scanners that can be used to dive into specific weaknesses. Third parties can also be employed to provide testing because sometimes, a

fresh set of eyes on the network can catch things the normal administrators would never think of. Third parties also often have more specialized and honed testing skills. Testing our networks is an ongoing process; the changes to both the network and the existing threats ensure the need for constant vigilance and monitoring.

When testing, don't forget to test your users. People forget what they don't constantly use, and security awareness is no different. Try out fake spam campaigns and see who clicks the "malicious" link, have rewards for catching people without appropriate badges, and have annual refresher courses. Train people to recognize a problem, and to know what response is appropriate to different problems [9].

3) Have a Budget to Implement Improvements

As you learn from events, yours or other organizations', you will likely find changes that need to be made. Security is not a product; security is a process. It is ongoing and takes time and resources. This may require the purchase of a new piece of electronic hardware (e.g. a firewall), or new software; (e.g. an intrusion detection system), physical hardware (e.g. a taller fence), or training time and instructors for your employees. When putting together your annual budget, there should be a line in there for those improvements that you find to be necessary throughout the learning process. A budget gives us the ability to fix things now and not to wait till next year. Next year may be much too late, and fixing a vulnerability is almost always less expensive than recovering from being hacked.

H. Be Aware of Your Public Information

Have you ever Googled yourself, a friend, or a co-worker? Chances are, you found a good deal of information on that person, ranging from phone numbers and addresses to personal data on social media. A control system can be treated the same way, and those data sources can provide a malicious user with a range of data that are invaluable when performing any sort of social engineering.

1) What Is Publicly Available About Your Network?

Both your control system and enterprise networks may be visible to the internet. Details about the internal structure, security, or devices should be kept private. It is not advisable to base your security on obscurity, but making things harder for the malicious user is always a good idea. Don't be the low hanging fruit!

2) What Is Publicly Available About Your Employees (Especially Leadership)?

Humans are targets for phishing attacks, and publicly available information about your employees can give the phishing attempt information enough to appear authentic. For example, in the HBGary attack several years ago, the attackers used information about the CEO to get a network administrator to open an unusual port and give out the root password [10].

III. Conclusion

Good security involves people, hardware, software, policy, and procedure, regardless of whether we are considering an enterprise network or a control system. However, the stakes are usually much higher with control system; it is unlikely that an enterprise network failure will endanger lives, but it is a frighteningly real possibility in a control system. If usability and efficiency demand that the control system be available on the Internet, the best practices laid out here provide a solid, secure base for a control system network to work from, and they also describe how to continue ensuring that your network is secure in the future.

By knowing, limiting, and monitoring access to the control system and implementing good defense in depth, we can create a control system that encourages security. Patching, updating and maintaining, ensuring physical security, and being aware of our public image help us to keep the level of security high. By monitoring the control system, we can catch the problems when they do happen, use our contingency plans to limit the effects, and then learn from those events to continuously improve. Security is an ongoing process.

IV. References

[1] K. Stouffer, J. Falco, and K. Scarfone, "Guide to Industrial Control Systems (ICS) Security," NIST Special Publication 800-82, June 2011.

[2] D. Watts, "Security and Vulnerability in Electric Power Systems," proceedings of the 35th North American Power Symposium, Rolla, MO, October 2003.

[3] P. Oman and M. Phillips, "Intrusion Detection and Event Monitoring in SCADA Networks," Critical Infrastructure Protection. Springer, Boston, MA, 2008, pp. 161–173.

[4] K. Barnes and B. Johnson, "Introduction to SCADA Protection and Vulnerabilities," Idaho National Engineering and Environmental Laboratory, March 2004. Available: http://www.inl.gov/technicalpublications/Documents/3310860.pdf.

[5] P. Oman, E. O. Schweitzer, III, and D. Frincke, "Concerns About Intrusions Into Remotely Accessible Substation Controllers and SCADA Systems," proceedings of the 27th Annual Western Protective Relay Conference, Spokane, WA, October 2000.

[6] United States Department of Homeland Security, "Recommended Practice: Improving Industrial Control Systems Cybersecurity With Defense-in-Depth Strategies," October 2009.

[7] P. Oman, E. O. Schweitzer, III, and J. Roberts, "Safeguarding IEDs, Substations, and SCADA Systems Against Electronic Intrusions," proceedings of the 3rd Annual Western Power Delivery Automation Conference, Spokane, WA, April 2001.

[8] E. O. Schweitzer, III, D. Whitehead, A. Risley, and R. Smith, "How Would We Know?" Proceedings of the 64th Annual Conference for Protective Relay Engineers, College Station, TX, April 2011.

[9] North American Electric Reliability Corporation, "CIP Standards," May 2006. Available: http://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx.

[10] P. Bright, "Anonymous Speaks: The Inside Story of the HBGary Hack," *Ars Technica*, February 2011. Available: http://arstechnica.com.

Previously presented at Resilience Week 2016, Chicago, IL, August 2016.