

Consider the consequences: A powerful approach for reducing ICS cyber risk

Richard Wyman

Received (in revised form): 11th January, 2017

Professional Control Systems Engineer, Idaho National Laboratory, P.O.Box 1625, Idaho Falls, ID 83415-3545, USA
Tel: +1 208-526-1249; E-mail: richard.wyman@inl.gov

Richard Wyman is a senior control systems engineer at Idaho National Laboratory (INL). During the last eight years, he has supported the United States Department of Homeland Security (DHS) Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) assessment and training programmes. As one of the original members of the ICS-CERT assessment team, Richard has evaluated over 100 control systems. Before his INL career, he worked as a project manager and technical lead for a northern California water utility, where he was responsible for the design and installation of a large distributed supervisory control and data acquisition (SCADA) system. In addition to his technical expertise, Richard has given presentations on controls and communications at several professional conferences and presented workshops on industrial communications, instrumentation and control systems in North America and Europe. He graduated with a Bachelor's Degree from Brigham Young University and a Master's Degree from the University of Washington in mechanical engineering.

ABSTRACT

Securing industrial control systems (ICS) or, for that matter, information technology (IT) systems is a never-ending battle. Cybersecurity subject matter experts (SMEs) secure their systems with the latest technology and threat actors develop new techniques to bypass these controls

in a constant arms race of attack and defend, attack and defend. This single-minded focus on responding to the latest threat often causes ICS cybersecurity SMEs to forget what they are defending, which is controlling and protecting the process. To be more specific, the cyber protections should prevent a threat actor from issuing malicious control commands and/or ensuring that a threat actor does not stop legitimate commands from reaching its objective. Unauthorised commands and the inability to issue commands have caused several high-profile impacts that resulted in significant damage in physical systems. This paper explores the relationship between cyber and physical systems by introducing a reference model that explains the cascading nature of impacts. While a cyberattack on an ICS originates in the cyber domain the most serious impacts occur in the physical domain. By understanding this concept, cybersecurity SMEs can make more targeted defensive measures in the cyber domain and add protections in the physical domain to significantly reduce ICS cyber risk.

Keywords: ICS cybersecurity, cyber-attacks, cyber/physical impacts, ICS Cyber Kill Chain, protection layers, risk analysis

INTRODUCTION

How does a cyber event on an industrial control system (ICS) impact a physical



Richard Wyman

process? On the surface, the answer to this question seems quite simple. A threat actor exploits vulnerabilities in an ICS and maliciously manipulates the system to cause physical consequences. In reality, the explanation is far more complex, but for now, breaking this simple explanation into sections and drilling deeper provides additional questions that if answered, will lead to a better understanding on how attackers exploit control systems. For example:

- ‘A *threat* actor exploits’ — Who are the threat actors? What are their motives? Curios? Ego? Financial? Vandalism? Terrorism? War? What are their capabilities? What are their resources? How do they exploit vulnerabilities? What are their techniques, tactics and procedures (TTP)? What is their level of access to the control system? Insiders? What are their signatures? How can they be monitored?
- ‘*vulnerabilities* in an ICS’ — Where are these vulnerabilities? Firmware? Operating system? ICS application? Network? Human? Can the information technology (IT)/ICS boundary be exploited? How secure is remote access? Who maintains the system? What is their level of expertise in securing control systems? Are the physical protections sufficient to prevent malicious tampering of the control devices, servers, and networking equipment? Are passwords easily guessed? These questions are a fraction of the issues that cybersecurity subject matter experts (SMEs) need to explore to develop a complete picture of the controls system’s vulnerabilities. For additional details in identifying vulnerabilities, refer to the ICS cybersecurity standards NIST 800-SP-82¹ and/or IEC 62443.²
- ‘and maliciously manipulates the system to cause physical *consequences*’ — This takes us back to the original question:

‘How does a cyber event on an ICS impact a physical process?’ which is the purpose of this paper.

By identifying and analysing threats, vulnerabilities and impacts, asset owners can develop a risk profile for their ICS. This information is essential in developing, prioritising and implementing mitigations to reduce ICS cyber risk. There are many excellent resources,³ including software tools like the Cyber Security Evaluation Tool (CSET™)⁴ that can help organisations identify and document threats and vulnerabilities in their control systems. Nevertheless, ICS cybersecurity consequences tend to be an enigma for most organisations. Yet understanding these impacts is a crucial component in identifying and mitigating cyber risks. Without this knowledge, asset owners are left to speculation and conjecture, which leads to fear, uncertainty, and doubt (FUD), or worse — complacency.

To help unravel this enigma, the organisation needs to understand there really is no single cyber/physical impact, but rather a series of cascading impacts. Figure 1 shows a generalised model of these ICS impacts, triggered by a cyber event that affects the CIA triad (ie confidentiality, integrity, availability). The loss of integrity and/or availability initiates a sequence of events that ripples through the cyber/physical layer to the physical layer where it can create a host of issues including equipment damage, loss of production and quality, as well as health, safety and environmental (HSE) impacts. More than likely it will also cause business impacts, and depending on the severity of the event, it could also generate community impacts.

Since the cascading impacts affect the control system, as well as its operations, maintenance, process, safety, product quality, and business, the organisation will

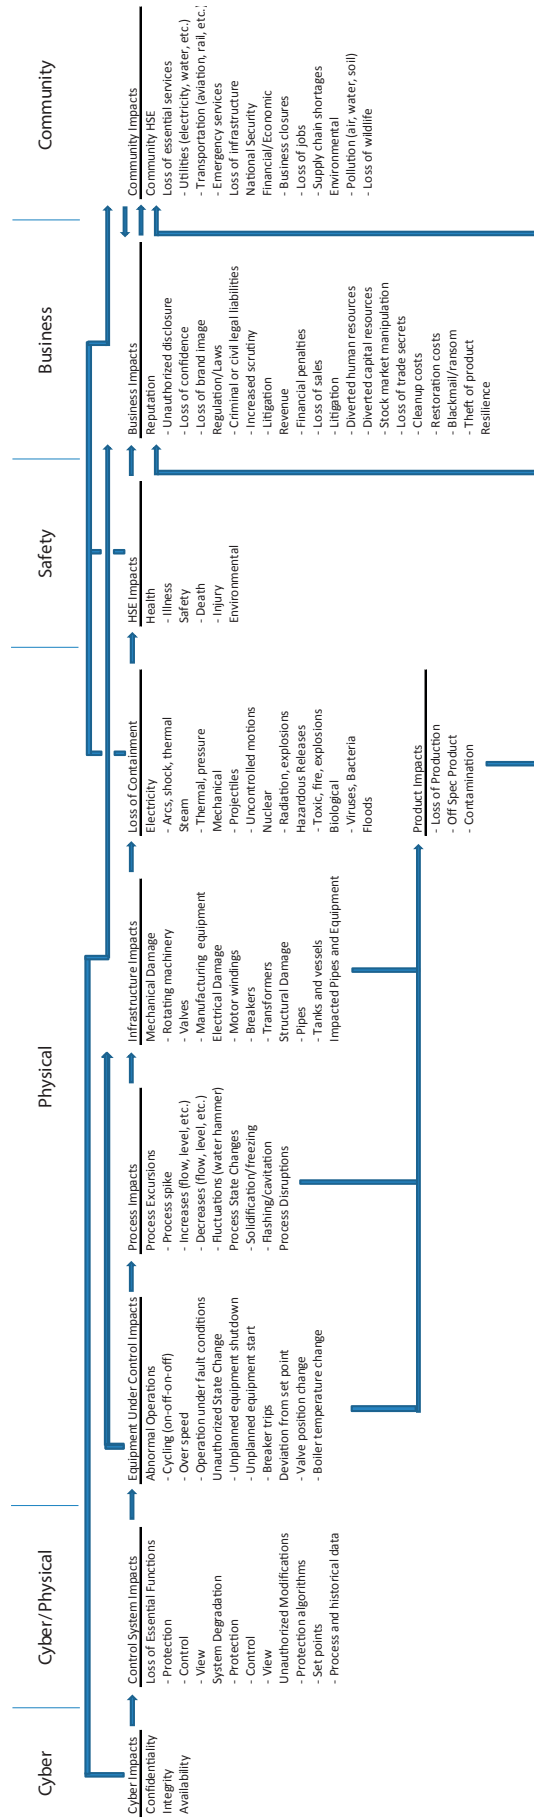


Figure 1 Cascading Impact Model

Source: created by Richard Wyman

need a team of experts to fully understand the risk of these cascading impacts. This team may include cybersecurity experts, control system engineers, operators, maintenance technicians, process engineers, safety engineers, quality control engineers, risk analysts and business managers. This model is generic. Domain experts may need to modify and adapt the model to accommodate specific processes or industry sectors. This paper will use the model to help explain the relationship between cyber and physical impacts.

By contrast, vulnerabilities and consequences in IT systems are tightly coupled and easier to comprehend. An exploited vulnerability that impacts the CIA triad leads directly to consequences, which may include information disclosure, system disruption and/or data alteration. These impacts occur within the cyber domain, although there may be business and community impacts that the organisation will need to address.

To bridge the gap between the cyber and physical domains, and gain a better understanding of the cascading cyber/physical impacts, this paper will:

- discuss the safety perspective of physical impacts,
- describe the ICS Cyber Kill Chain model and Consequence-driven Cyber-Informed Engineering (CCE),
- explain the cyber/physical interface,
- discuss protection layers,
- examine the path of a control command to physical impact, and
- analyse the impact of cascading events.

SAFETY PERSPECTIVE OF PHYSICAL IMPACTS

Most organisations have strong safety cultures, especially those that operate hazardous processes, such as nuclear facilities, refineries, and chemical plants that

transport, store, and process hazardous materials. Because of their potential for generating high-consequence events (eg toxic releases, fire, explosions), these industries are heavily regulated with the goal to reduce risk. Many of the regulations stipulate change control, asset identification, training, physical controls, policies and procedures, audits, etc.⁵ — concepts that are familiar to most cybersecurity professionals. The regulations also require a safety assessment or process hazard analysis (PHA). Examples of PHAs include ‘What-If’, ‘Checklist’, ‘Hazard and Operability Study (HAZOP)’, ‘Failure Mode and Effects Analysis (FMEA)’ and ‘Fault Tree Analysis’.⁶ These are structured and mature methodologies designed to identify human, operational, equipment and process problems that could cause HSE impacts.

Safety engineers use PHAs to reduce safety risk, not cybersecurity risk. Originally, control systems and safety instrumented systems (SISs), responsible for mitigating safety hazards, were installed in isolated environments. As a result, safety engineers and control system engineers did not view cyberthreats as a risk they needed to mitigate. This is starting to change. Several influential ICS cybersecurity SMEs are calling for action in doing more to secure cyber systems that protect plant and process safety.

Cusiano and Gruhn acknowledge that control and safety systems ‘are vulnerable to cybersecurity breaches’ and emphasise that ‘traditional’ hazard mitigations, responsible for protecting the process, do not take into account the cyberthreat. They go on to make the case that additions to process safety and cybersecurity standards will help close this safety/security gap by requiring cyber vulnerability and risk assessments for safety systems.⁷

Anderson and Price made similar points on the vulnerabilities of control systems

and issued a challenge for developing new methodologies based on ‘Cyber-Informed Engineering’, in a paper given at the ‘International Conference on Computer Security in a Nuclear World: Expert Discussion and Exchange’ in June 2015. The following excerpt from their abstract emphasises these points:

‘Current methodologies focus on equipment failures or human error as initiating events for a hazard, while cyber-attacks use the functionality of a trusted system to perform operations outside of the intended design and without the operator’s knowledge. These threats can bypass or manipulate traditionally engineered safety barriers and present false information, invalidating the fundamental basis of a safety analysis. Cyber-threats must be fundamentally analysed from a completely new perspective where neither equipment nor human operation can be fully trusted. A new risk analysis and design methodology needs to be developed to address this rapidly evolving threatscape.’⁸

The convergence of safety and security highlights an important point made in NIST 800–SP–82:

‘Safety assessments are concerned primarily with the physical world. Information security risk assessments primarily look at the digital world. However, in an ICS environment, the physical and the digital are intertwined and significant overlap may occur.’⁹

As pointed out earlier, safety by no means is the exclusive domain of the chemical and nuclear sector. The public holds the energy, transportation, dams, food and agriculture, water and wastewater, as well as other sectors, to high safety standards. No one wants to drink contaminated water, fly on planes that are not regulated or live downstream of a dam that is not safe. Because the safety risks associated

with each sector varies, the methodologies and rigour of sector specific safety assessments and regulations also vary. Understanding these sector specific safety assessments and the regulatory environment can help identify potential impacts during a cybersecurity assessment. Other resources for identifying impacts for security assessments include safety inspectors, engineers, operators, maintenance technicians, risk analysts, consultants, trade associations, academia, national laboratories, vendors and regulators.

Turning again to NIST 800–SP–82, the standard stresses the importance of communications between security and safety stakeholders:

‘The personnel responsible for the information security risk assessment must be able to identify and communicate identified risks that could have safety implications. Conversely, the personnel charged with safety assessments must be familiar with the potential physical impacts and their likelihood developed by the information security risk assessment process.’¹⁰

This is an excellent recommendation that, if followed, will lower both safety and security risk.

THE ICS CYBER KILL CHAIN AND CONSEQUENCE-DRIVEN CYBER-INFORMED ENGINEERING (CCE)

In the movies, the hero gives a quick few taps on a computer keyboard and at the last minute, the nuclear warhead hurtling towards a major city self-destructs. As those involved with cybersecurity know, this movie scenario is no more realistic than the simple explanation given at the beginning of this paper on how an attacker exploits an ICS. The reality is that high-impact attacks take planning, time and persistence.

In their paper, ‘The Industrial Control System Cyber Kill Chain’,¹¹ Assante and Lee give a more realistic view on how attackers exploit ICS. While the ICS Cyber Kill Chain is based on Lockheed Martin’s popular Cyber Kill Chain model, the authors have modified it to represent the unique characteristics of control systems. The ICS Cyber Kill Chain describes two stages for compromising a system. The first stage is similar to Lockheed Martin’s Cyber Kill Chain. The attack pattern follows the workflow for a typical IT breach where the impacts, with the exception of the business and community layer, are confined to the cyber layer. The purpose of this stage is to gather information on the control system, networks, vulnerabilities, and processes. By now, the threat actor has most likely defined goals and objectives for affecting the physical process, harming the business, and possibly the community.

The attacker uses this information during the second stage to develop, test, deliver, install/modify, and execute an ICS attack that creates the desired physical consequences. This is the stage where the attacker compromises the integrity or availability of the cyber system to trigger the cascading impacts with the ultimate goal of creating production problems, damage equipment, and/or cause HSE impacts. Threat actors have done this with devastating effect. For example:

- In 2007, researchers at Idaho National Laboratory (INL) demonstrated they could destroy a 2.25 MW generator through a cyberattack.
- The Stuxnet virus, discovered in 2010, destroyed up to 1000 centrifuges for processing enriched uranium at the Natanz fuel enrichment plant in Iran.
- On 23rd December, 2015, a highly trained and well-funded organisation demonstrated that they could

cause widespread power outages in the Ukraine by tripping breakers and performing other malicious activities. These outages lasted up to six hours and affected close to 225,000 people.¹²

Each of these cyberattacks took planning, skill, and a strong understanding of the process. It is important to learn from these attacks and understand the methodologies discussed in the ICS Cyber Kill Chain. In addition, the INL Mission Support Center produced a concept paper, ‘Consequence-driven Cyber-informed Engineering (CCE)’, which provides additional understanding of adversary activity including ‘top-tier, highly resourced adversaries’.¹³ The ICS Cyber Kill Chain and CCE provides valuable insight on how to protect the control system and ultimately the process. As an example, knowing an attacker needs information on the process for developing and executing an attack that produces physical impacts should increase the organisation’s efforts for protecting and securing this information. This would include, but not be limited to, process/piping and instrumentation diagrams (P&ID), process flow diagrams, loop diagrams, instrument lists, and operation manuals. Protecting these documents should be as important as securing ICS network diagrams, ICS system descriptions, data flow diagrams and IP address listings. It also means the organisation should protect their control systems from data exfiltration, especially the ICS configuration database, historical database, controller configuration files, and other information a threat actor could use to understand and develop a physical event. Confidentiality is typically not a high priority in establishing ICS security controls, but cybersecurity SMEs should seriously consider confidentiality mitigations if supported by a risk assessment.

THE CYBER/PHYSICAL INTERFACE

ICSs are referred to as cyber/physical systems (CPSs) for a reason. They rely on a cyber infrastructure for collecting data, making control decisions, providing an interface for the operator, storing information, issuing alarms, and many other uses. The main purpose of these systems, however, is to change the state of equipment under control (EUC), which resides in the physical domain. The infrastructure for performing these activities ranges from a simple controller for turning on a lawn sprinkler to countless computers, networks and controllers for refining oil or processing wastewater. Since the control system spans both the physical and cyber domains, it is important to understand the interface between these two domains.

Control system engineers commonly use the Purdue reference model, as shown in Figure 2, to describe relationships between ICS subsystems and the enterprise network, but this model can also help explain the cyber/physical interface. Level 0, the process layer, includes EUC (eg motors, valves, actuators, fans) and sensors (eg pressures, flows, levels, temperatures, alarms, status). Level 1, referred to as basic control, includes controllers such as programmable logic controllers (PLCs), remote terminal units (RTUs), intelligent electronic devices (IEDs) and distributed controllers. SISs are an independent system




Figure 2 *Purdue Reference Model*

Source: SANS, Luciana Obregon¹⁴

that will be covered later. Irrespective of what you call them, controllers do three important things — provide input/output (I/O), control, and communication with Level 3 supervisory control components.

The I/O in the controllers bridge the physical and cyber domains. The two most common types of I/O are discrete and analog. Discrete signals are binary in nature; that is, they are either on or off, open or closed, in alarm or normal. In most cases, discrete input signals are nothing more than a process change that causes a switch to complete or break a circuit. This circuit is monitored by the discrete input channel on the controller, which tells the controller the state of the input (eg the pump is running or the high-level switch is in alarm). The discrete output on the controller activates or deactivates an internal mechanical or electronic relay that makes or breaks an electrical circuit. This circuit will turn the EUC on or off depending if the circuit is energised or de-energised.

Analog signals are a little more complicated. The analog transmitter (eg pressure, flow, etc.) generates a variable analog signal (eg 4 to 20mA, 1 to 5 volts, etc.) in proportion to the value of the process variable. The transmitter sends the signal to the analog input of the controller where an analog to digital converter (ADC) changes the signal to a digital signal (eg bits and bytes). Conversely, the output in the controller takes the bits and bytes and converts them into a variable analog signal using a digital to analog (DAC) converter. This signal tells the EUC what to do, such as open the valve to 75 per cent of range.

There are several observations to make regarding the cyber/physical interface:

- The control system funnels all commands that have a physical impact, good or bad, through the controllers' output.
- Although there are numerous techniques

a threat actor can use for influencing the physical domain, they include either sending unauthorised control commands to the controller's output or preventing legitimate commands from reaching the controller's output. A small sampling of these techniques include:

- hijacking an operator workstation to issue unauthorised control commands
- intercepting and modifying control commands while in transit over the network
- spoofing the process/input data to trick an operator into issuing bad control commands
- vandalising the control system or creating a denial of service to prevent an operator from issuing control commands.
- The physical domain is isolated from the cyber system at the I/O. An observer on the physical control loop cannot tell how complex or simple the cyber portion of the control system is by watching control commands issued by the controller. In other words, the observer, viewing the control loop, cannot tell if the refinery pump motor is controlled by a simple lawn controller or countless computers, networks, and distributed controllers, or for that matter, an operator manually closing a switch.
- The cyber domain is partially isolated from the physical domain at the I/O. An observer on the cyber domain may glean context on the process by viewing process graphic displays, reading point descriptions, and examining the programs for the controller, but they cannot tell how the engineers designed the physical control loop.

PROTECTION LAYERS

An operator once quipped he doubted an attacker could do any more damage to the system than he or his colleagues have done.

People make mistakes. Operators issue commands at the wrong time or fail to issue commands when urgently needed. When an operator makes a mistake there may be process disruptions, a tank may over-fill requiring a costly cleanup, or workers may need to discard a product batch, but generally, there are no Armageddon melt-downs with these types of mistake. Bad 'stuff' happens all the time, but operators are usually able to deal with the resulting problems. When the operator confessed that he and his colleagues made mistakes, he unknowingly highlighted the importance of protection or safety layers. When there is a bad incident, the investigation usually shows there was a breakdown in one or more protection layers.

Protection layers are the safety version of defence-in-depth (DID). Safety engineers, when conducting a PHA, often use a risk assessment method called layers of protection analysis (LOPA) to reduce safety hazards. LOPA, as shown in Figure 3, explains how the safety and cyber domains work together to protect people, equipment, and the environment, but it also shows how a threat actor can exploit the cyber domain for nefarious purposes. Additionally, LOPA illustrates how physical



Figure 3 Layers of Protection Analysis

Source: http://www.safety-s2s.eu/modules.php?name=s2s_wp4&idpart=2&op=v&idp=750¹⁵

controls, external to the cyber system, not only reduces safety risk, but also reduces security risk from a cyber exploit.

Referring to Figure 3, ICS cyber-security SMEs and control systems engineers can discover important insights by viewing each LOPA layer from a security perspective:

- *Process/Process Design* — Some processes are more risky to HSE than other processes. For example, just over 100 years ago, water utilities started using chlorine gas to disinfect water. Chlorine gas is a toxic substance that presents a hazardous risk to its employees and the surrounding community; as a result, water utilities are migrating to less hazardous forms of disinfectant. This migration not only removes the safety risk, but also eliminates the security risk of maliciously operating the process to cause a chlorine release, if that were even possible.

Evaluating EUC for cyber impacts may be another opportunity to lower security risk. For instance, how does the malicious operation of EUC affect its reliability and ability to control the process? For example, some motors may be more robust in handling abusive commands from a threat actor than other motors.

- *Basic Process Control System (BPCS)* — As mentioned previously, one of the key functions of controllers is to execute control logic. It does this by monitoring process inputs and issuing commands to the EUC to bring the process back into specification. The controllers also protect the EUC from damage. Interlocks are the primary mechanism for accomplishing this objective. As an example, a low suction pressure can damage a pump. To prevent this, control system engineers program an interlock that shuts down the pump if the pressure drops below a given threshold. The

interlock can also be non-cyber, such as a low-pressure switch that shuts off the pump on low suction pressures.

Although non-cyber interlocks may be more expensive to install and less flexible to modify, they have one huge advantage over cyber interlocks. A threat actor cannot compromise a non-cyber interlock with a cyberattack. In their paper, ‘The Case for Simplicity in Energy Infrastructure’, Assante, Roxey and Bochman point out that:

‘Although these modern technology enhancements will result in greater productivity and efficiency contributing trillions of new value to the global economy, we are also unlocking an equally powerful dark side that can negate these advantages.’¹⁶

While the authors are not advocating a wholesale return to non-cyber control, they are recommending that we ‘re-engineer selected elements of the grid’ and ‘use analog, nondigital, or purpose-built digital circuits’ where justified by a risk assessment.¹⁷

NIST 800–SP–82 further highlights the benefits of using analog control systems as a means ‘to prevent the physical process from entering an undesired state in situations when the digital control system is unavailable or corrupted. Analog controls include regulators, governors, and electromechanical relays.’¹⁸

- *Alarm, Operator Action* — Referring back to the Purdue model in Figure 1, the controllers in Level 1 send real-time data to Level 2, the supervisory control layer. From this layer, operators acknowledge alarms, analyse historical trends, and monitor real-time displays. If the controller in Level 1 does not keep the process in check, depending on the severity of the situation, the operator has the option to continue to monitor

the process, manually intervene if necessary, or initiate emergency operations. This only works if the supervisory layer is available and has integrity, that is, an attacker has not compromised the control system.

Measures to ensure availability, such as redundant networks, redundant servers and backup control centres, were designed to mitigate equipment failures, not cyberattacks. Redundant systems may not provide availability if the cyberattack also compromised the backup system. Attacks that negate integrity also present serious challenges. An attacker may be able to deceive an operator into thinking everything is okay if they disable alarms and spoof the data in the operator displays to show normal operations.

There are mitigations, such as an independent analog alarm system, non-digital gauges (eg pressure, level, flow, etc.), manual control valves and physical breakers. These measures, however, can be costly, can substantially increase operator intervention, and can be difficult to maintain.¹⁹

- *Safety Instrumented System (SIS)* — SISs are cyber-based systems designed to bring the plant into a safe state if the BPCS fails to keep the process within operating boundaries. These systems are fault-tolerant and should be independent of the BPCS, meaning that they have separate sensors, valves, and logic solvers or controllers. Given the emerging cyberthreat, Cusimano and Gruhn cite two problems with how organisations design, install, and maintain these systems: (1) there is often some level of integration between the SIS and the BPCS, including shared network access, and (2) traditional PHAs do not consider the impacts from a cyberattack. These are serious problems given the tight coupling between process safety and cybersecurity. An attacker can potentially disable, or even worse,

exploit all three cyber layers of protection (BPCS, Alarms, and SIS) with a single attack.²⁰

Usually, only hazardous processes like nuclear power plants, hazardous chemical plants and refineries use SISs. Less hazardous operations rely on control systems or non-cyber protections to safely shut down EUC.

- *Active Protection* — Examples of active protection devices include relief valves, rupture disks, breakers, motor control protection, and other measures that step in to mitigate the effects of an uncontrolled process. These controls will mitigate the effects of the uncontrolled process, regardless if a safety or security event triggered the incident.
- *Passive Protection* — These measures contain or mitigate the impacts of a safety event through passive mechanisms, but they also mitigate cyber-triggered events. Examples include dikes, shields on rotating equipment, nuclear containment structures and reservoir spillways.
- *Emergency Response* — Includes actions to limit damage, injury, or the loss of life. These activities include evacuation, firefighting, incident response, public notification, etc.

This layered approach for preventing safety impacts explains why human error by an operator is relatively benign. The protection layers, especially the cyber layers (eg BPCS, Alarms, SIS), should prevent an operator's actions from causing a serious mishap. If a threat actor were able to disable one or more of the protection layers in a cyber system, however, the impacts of a mistake could be more serious. In many cases, an attacker must disable or bypass these protection layers to generate high-impact events.

The LOPA model works well for identifying security risks associated with safety issues; however, the model is not

appropriate for analysing other types of cyber impacts, such as loss of product quality or loss of production.

FROM CONTROL COMMAND TO PHYSICAL IMPACT

In traditional control system architectures, Level 1 controllers in the Purdue model execute control logic and communicate with the Level 2 Supervisory Control layer. Typically, a PLC, RTU, or IED will automatically generate control commands (Level 1 Basic Control) in response to process inputs or execute Supervisory Control commands issued by an operator from the HMI (Level 2 Supervisory Control). Controller-generated commands are more difficult to exploit since the control logic and corresponding process

inputs used to generate the control outputs typically reside in the same controller. As a result, the process inputs, control logic, and outputs are not directly exposed to network attacks. By contrast, the attack surface for supervisory issued commands is much greater. These commands traverse networks where an attacker can intercept and modify these supervisory commands. A compromised node can also masquerade as a legitimate host since most control system protocols do not authenticate or authorise control actions, giving an attacker the opportunity to issue unauthorised commands. System complexity can also increase the attack surface, giving a threat actor more options for exploiting the control system. Simple ICS architectures are easier to defend than complex system architectures.

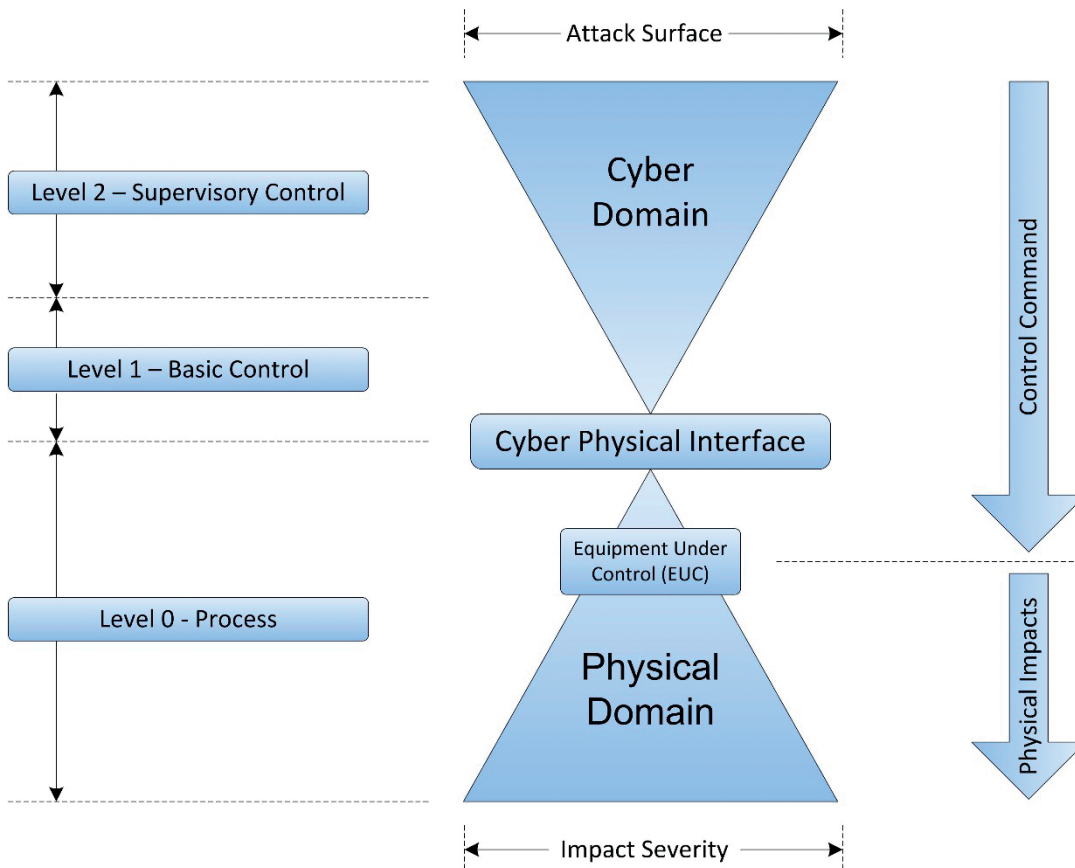


Figure 4 Cyber and Physical Domains Diagram

Source: created by Richard Wyman

As a rule, accessing and exploiting the supervisory level is easier than accessing and exploiting the basic control level. Figure 4 shows the attack surface shrinking, as the command gets closer to the cyber/physical interface. To emphasise this point, Basic Control is independent of Supervisory Control; however, Supervisory Control is dependent on Basic Control. In most cases, Basic Control will continue to generate controller commands if an event or disaster compromised or destroyed the Supervisory Control; however, operators will not be able to issue Supervisory Control commands if the Basic Control is not available.

Some of the newer control systems, currently deployed by industry, do not follow traditional ICS architectures. Because the data flows and corresponding attack surfaces are different, the model shown in Figure 4 will not apply to these newer systems.

This brings up the importance of protecting Basic Control from cyberattacks. Typically, the primary and sometimes only protective logic for preventing equipment damage runs in the Level 1 controller. An attacker will target these defences by bypassing or disabling these protections. When justified by a risk assessment, engineers should design controls at the cyber physical interface to limit potential consequences. In addition, maintenance technicians and/or control system engineers should secure backup copies of controller configuration files used to restore the controllers, making it more difficult for a threat actor to tamper with controller programs. Physically protecting controllers is also crucial. The organisation should secure controllers in locked cabinets and monitor access to the controller with intrusion detection alarms.

Referring again to Figure 4, just as the attack surface expands moving away from the cyber/physical interface in the

cyber domain, so does the severity of hazards in the physical domain. Initially the cascading impacts are serial and easier to contain; however, as they continue to migrate unchecked they branch out as parallel impacts, making it much more costly and difficult, if not impossible, to mitigate. As will be shown in the next section, it is easier to disrupt a control loop that is cycling from an energised to a non-energised state than it is to stop the potential HSE impacts and production impacts from a pump failure.

ANALYSING THE IMPACT OF CASCADING EVENTS

An organisation can use the model (refer to Figure 1) to analyse impacts associated with the sequence of events triggered by a cyber-attack. As an example, engineers want to evaluate the impacts of maliciously cycling the controls for a water distribution pump motor. By developing a list of questions for researching the relationships and interactions between cascading events, analysing the physical consequences of each event, and evaluating existing measures to prevent the events from propagating, a deeper understanding can be developed for the effects caused by a given cyber-attack. The engineers can use these insights to develop mitigations and/or protection layers for reducing or eliminating the potential risks. Their objective is to stop the cascading events before they create serious consequences. Table 1 shows an example on how analysts can use the cascading impact model to evaluate the risks of an attacker maliciously cycling the controls of a water distribution pump motor.

The MitM attack, used in this example to issue the malicious commands, is like knocking over the first tile in a row of dominos. Just as the fall of the first domino starts a spectacle of toppling dominos,

Table 1: Cascading Impact Analysis

<i>Layer</i>	<i>Potential Event</i>	<i>Impact Analysis</i>
Cyber		
<u>Cyber Impacts</u>		
Loss of integrity	Attacker executes a Man in the Middle (MitM) attack between the polling host and pumping plant PLC.	What is the likelihood of executing a MitM attack? What are the mitigations in place to prevent a MitM attack? What are some of the other measures an attacker could take to compromise system integrity to issue cycling commands to the PLC? Can an attacker impact multiple pumps at the same time?
Cyber/Physical		
<u>Control System Impacts</u>		
Unauthorised modification of set point	Attacker issues commands to cycle the PLC output for controlling a pump motor.	What are the measures in place to prevent and/or mitigate cycling commands? Protection logic in the PLC (for example, time-outs between pump stops and starts)? Can an attacker bypass the protection logic? Does the PLC protect the interlocks and protection logic from unauthorised modifications?
Physical		
<u>EUC Impacts</u>		
Abnormal control (cycling)	The PLC issues a cycling discrete signal to the pump motor by energising and de-energising the control circuit.	Are there physical controls to prevent or disrupt cycling of energising and de-energising the control signals?
<u>Process Impacts</u>		
Process excursions (pressure fluctuations)	Cycling pump starts and stops, creating pressure fluctuations in the pipe (eg fluid hammer).	Can starting and stopping pumps cause a water hammer? At what point does it cause operational impacts? What are these operational impacts? Can they impact customers? Create noise? Break pipes?
<u>Infrastructure Impacts</u>		
Electrical damage (motor windings)	The cycling inrush current heats motor windings.	Are there physical protections such as breakers or thermal overloads to prevent overheating the windings? What is the tolerance of the pump windings to resist shorts and/or break down of the insulation? Are there other motors more resilient to cycling commands? How long will it take to fix or replace a damaged motor?
<u>Infrastructure Impacts</u>		
Structural damage (pipes)	The cycling creates water hammers, which breaks the pipes.	At what point do the pipes break? Where do they break? Are there physical devices such as surge tanks, dampers, or relief valves to prevent breaks from pressure surges? How long will it take to repair the pipe?
<u>Loss of Containment</u>		
Electricity (arcing, thermal)	The cycling inrush currents cause arcing and elevated temperatures.	What are the protective measures in place to contain arcs? Can the thermal increases cause fire? If so, what measures are in place to prevent, detect or mitigate fires?

<i>Layer</i>	<i>Potential Event</i>	<i>Impact Analysis</i>
<u>Loss of Containment</u> Floods	Broken pipe causes flooding.	What structural or environmental damage can a flood cause? What about homes or roads in the surrounding area? Can the floodwaters be safely diverted?
<u>Product Impacts</u> Loss of production	Damaged pump and pipe prevent delivery of water to customers.	Are there other pumps and/or pumping plants to continue water distribution? Are there alternative pipe routings to deliver water? How long will it take to restore service? Are portable pumps a viable option for maintaining water deliveries?
<u>Product Impacts</u> Contamination	Broken pipe exposes potable water to containments.	Is water infiltration from ground water a possible source of water contamination? How are contaminants detected? How are they mitigated? How is the public notified of possible contamination?
Safety <u>HSE Impacts</u> Health, Safety, Environmental	Arcing, fire, floods, and water contamination are potential impacts that could affect HSE.	What is the likelihood of HSE impacts due to other impacts? What are these potential HSE impacts? Can they be eliminated or mitigated? Are staff trained to mitigate HSE impacts? How does the organisation interface with emergency services?
Business <u>Business Impacts</u> Reputation (loss of brand image, loss of confidence)	Customers and government regulators are understandably concerned about the organisation's ability to safely deliver water when needed.	What information should the organisation release on the incident? How will they notify interested stakeholders on repair status? How does the organisation interface with law enforcement and other federal government agencies that coordinate a cyber response? How will they rebuild trust?
<u>Business Impacts</u> Revenue (loss of sales, litigation, cleanup costs, restoration costs)	The organisation incurs substantial costs to recover from a cyberattack.	Does the organisation have sufficient cash reserves to recover from a cyber event? Is financial risk transfer (insurance) an option? How are costs tracked and managed during an emergency? What are the potential losses due to decreased revenue during an outage?
Community <u>Community Impacts</u> Community HSE	Damage to utility infrastructure incurs HSE impacts.	Can damage to physical infrastructure caused by a cyberattack create HSE impacts? If so, how? Fire? Flood? Contamination? How substantial are these impacts?
<u>Community Impacts</u> Loss of essential services	Loss of water deliveries impacts other sectors.	What are the potential impacts to the utility's customers on loss of service? How long can businesses sustain a loss of water before incurring significant impacts? What mitigations can the utility provide to lessen the impact on its customers? How will firefighting be impacted?

an attacker can trigger the sequence of cascading impacts by forcing the PLC's control output on or off using a MitM attack. A MitM attack is not the only method an attacker can use to generate unauthorised cycling commands, however. These alternative attack scenarios could potentially produce the same devastating effects as a MitM attack, provided they bypassed the protection logic and delivered the same cycling commands directly to the cyber/physical interface. Assuming the control commands were the same signal as it passed through the cyber/physical interface, the analysis, downstream of the interface, should also be the same, irrespective of the triggering attack.

Conversely, an attacker can use a MitM attack to execute other types of malicious commands or stop legitimate commands from reaching the EUC. This will produce an entirely different impact scenario than the example outlined in Table 1, giving an attacker the flexibility to develop specific control commands to meet a specific impact objective. These objectives may include overflowing a water distribution tank, causing service outages, or bypassing low suction protection measures. As with the pump-cycling example, engineers can use the cascading impact model to evaluate the risks for each cascading impact that may lead to overtopping the tank or curtailing water deliveries.

The fact that a threat actor can produce the same impact using different types of attacks or generate different impacts using the same type of attack should be a concern to defenders of control systems. Although their ICS may be hardened to prevent attacks such as MitM, a threat actor can continue to probe the control system for other vulnerabilities they could exploit to issue rogue commands.

In addition, the permutation of attacks and impacts creates an unwieldy combination of scenarios to analyse; however, the

analyst can significantly reduce this list by focusing on impacts and not the associated attack that generated the impact. Referring to the domino analogy, the analyst does not need to know if the first domino was pushed with a finger, blown over with a fan, or hit with a spoon; only that it fell on the adjoining tile initiating the action of the toppling dominos. In a similar fashion, the analyst, looking strictly at impacts, does not need to examine the type of attack that triggered the impacts — only that it generated a specific control output to produce a specific set of cascading impacts.

Understanding the pattern of abusive control output is crucial to preventing its propagation. This analysis can be performed from the control system perspective (top down) or from the safety perspective (bottom up). The control system perspective is an iterative approach that determines the pattern of abusive commands capable of generating the impact under evaluation. The safety perspective chooses a specific impact for evaluation and works backward to identify the abusive control commands needed to generate that impact. Once the control system engineer has ascertained the abusive command pattern for the impact under analysis he/she can develop protective measures to block the abusive control signal from affecting the process. As an example, engineers can implement measures to prevent the cycling commands from reaching the EUC knowing that these commands may damage pumps.

CONCLUSION

ICS cybersecurity experts primarily focus on reducing system vulnerabilities to secure control systems, although organisations are starting to deploy threat monitoring in ICS networks to identify malicious activity. While these measures are essential as part of a comprehensive

approach for reducing ICS cyber risk, they are not sufficient for preventing a dedicated and well-funded foe from exploiting control systems to generate high-impact events. Control systems are increasing in complexity making them more difficult to secure. Threat actors use zero day attacks to exploit what were thought to be secure control systems. New attack techniques are able to bypass perimeter boundaries. Policies and procedure violations leave systems vulnerable to exploit. Clearly more needs to be done to globally secure ICSs.

Control systems have a significant advantage over IT systems. IT impacts are more tightly coupled with the exploit than control system impacts. An attacker who breaches the confidentiality of an IT system immediately generates an impact that exposes sensitive information that may include emails, pictures, drawings, etc. By contrast, successfully exploiting a control system vulnerability only triggers the cascading impacts. The cascading impact model, described Figure 1 of this paper, shows how the triggering event that originates in the cyber domain propagates to the physical domain where it causes physical damage. By understanding how these cascading impacts propagate, control system and safety engineers can develop cyber and physical protective layers to block malicious control commands from affecting the process.

Organisations often overlook or ignore cyber/physical impacts when performing an ICS cybersecurity risk assessment; however, they can significantly lower ICS cyber risk by including impacts in a comprehensive risk assessment. If measures implemented to reduce ICS vulnerabilities fail to thwart a cyber-attack, protective layers designed to stop cascading impacts may be the only defence for preventing a deadly attack that causes physical damage.

NOTES

- (1) For the full publication, cf. National Institute of Standards and Technology (NIST) Special Publication 800–82 Revision 2 (2011), ‘Guide to Industrial Control Systems Security’, available at <http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf> (accessed 28th November, 2016).
- (2) For more information, cf. the ISA99 Committee wiki page, available at <http://isa99.isa.org/ISA99%20Wiki/Home.aspx> (accessed 28th November, 2016).
- (3) For a list of resources, cf. the SCADAhacker homepage, available at <https://scadahacker.com/library/index.html> (accessed 28th November, 2016).
- (4) For more information, cf. the CSET webpage, available at <https://ics-cert.us-cert.gov/Downloading-and-Installing-CSET> (accessed 20th March, 2017).
- (5) For a list of regulations, cf. the OSHA webpage, available at https://www.osha.gov/pls/oshaweb/owadisp.show_document?p_table=STANDARDS&p_id=9760 (accessed 28th November, 2016).
- (6) *Ibid.*
- (7) For more information, cf. Cusimano, J. and Gruhn, P. (2016), ‘Integrating Industrial Control System (ICS) Cybersecurity with Process Safety Management’, available at <http://www.aesolns.com/wp-content/uploads/2016/04/Integrating-ICS-Cybersecurity-and-Process-Safety-Management.pdf> (accessed 28th November, 2016).
- (8) For more information, cf. Price, J. and Anderson, R. (2015), ‘Cyber-Informed Engineering: The Need for a New Risk Informed and Design Methodology’, available at <http://www.osti.gov/scitech/servlets/purl/1236850> (accessed 28th November, 2016).
- (9) NIST SP 800–82 Rev. 2, ref. 1 above.
- (10) *Ibid.*

- (11) For more information, cf. Assante, M. and Lee, R. (2015), 'The Industrial Control System Cyber Kill Chain', available at <https://www.sans.org/reading-room/whitepapers/ICS/industrial-control-system-cyber-kill-chain-36297> (accessed 28th November, 2016).
- (12) For more information, cf. SANS, 'Analysis of the Cyber Attack on the Ukrainian Power Grid' available at https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf (accessed 24th January, 2017).
- (13) For more information, cf. Freeman, S., St. Michel, C., Smith, R., and Assante, M., OTSI, 'Consequence-driven Cyber-informed Engineering (CCE)' (2016) available at <https://www.osti.gov/scitech/biblio/1341416-consequence-driven-cyber-informed-engineering-cce> (accessed 20th March, 2017).
- (14) For more information, cf. SANS, Luciana Obregon (2015), 'Secure Architecture for Industrial Control Systems', available at: <https://www.sans.org/reading-room/whitepapers/ICS/secure-architecture-industrial-control-systems-36327> (accessed 20th March, 2017).
- (15) For more information cf. 'S2S — A Gateway for Plant and Process Safety — Training LOPA' available at http://www.safety-s2s.eu/modules.php?name=s2s_wp4&idpart=2&op=v&idp=750 (accessed 20th March, 2017).
- (16) For more information, cf. Assante, M., Roxey, T. and Bochman, A. (2015), 'The Case for Simplicity in Energy Infrastructure', available at <https://www.csis.org/analysis/case-simplicity-energy-infrastructure> (accessed 28th November, 2016).
- (17) *Ibid.*
- (18) NIST SP 800-82 Rev. 2, ref. 1 above.
- (19) *Ibid.*
- (20) Cusimano, J. and Gruhn, P, ref. 7 above.