



SANS Institute

Information Security Reading Room

Collection and Analysis of Serial-Based Traffic in Critical Infrastructure Control Systems

Jonathan Baeckel

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Collection and Analysis of Serial-Based Traffic in Critical Infrastructure Control Systems

GIAC (GCIA) Gold Certification

Author: Jonathan Baeckel, baeckel@protonmail.com
Advisor: Bryan Simon

Accepted: January 15, 2021

Abstract

There is a blind spot the size of a 27-ton, 2.25-megawatt maritime diesel generator in the world's critical infrastructure control system (CICS) landscape. Compared to typical IT systems, CICSs are composed of a much larger ratio of non-routable traffic, such as serial-based Fieldbus communications, than their IT-based brethren, which almost exclusively rely on TCP/IP-based traffic. This traffic tells field devices to take actions and reports back process status to operators, engineers, and automated portions of the process. As vital as it is to the process, this specialized traffic is routinely ignored by Operational Technology (OT) architects and analysts charged with defending this type of system. They tend to favor a TCP/IP only approach to traffic collection and analysis that is more geared toward an IT-only environment. This paper analyzes Stuxnet to determine the effect that serial communication monitoring and analysis may have on the situational awareness of such an event. It will pose several questions. Could the attack have been detected without the availability of known Indicators of Compromise (IoC)? Would the attack have been detected sooner? Would there have been no effect at all? This information may help organizations pursue a risk-based approach to architecting a CICS traffic collection and analysis system.

1. Introduction

There are examples of catastrophic failures of physical equipment caused by the remote manipulation of processes, such as the Aurora Generator Experiment, which dramatically destroyed a 2.25 MW generator in under three minutes by remotely operating breakers. Stuxnet, which was able to destroy entire cascades of uranium hexafluoride gas centrifuges, is another example.

The purpose of this research is to briefly describe a high visibility attack on control systems that supports critical infrastructure and examine how collecting serial-based traffic for analysis may have affected the situational awareness of those supporting the control system process. The research explores the current state of traffic monitoring in Industrial Control Systems (ICS), discusses the importance of situational awareness, how serial collection might affect situational awareness within control systems, and it covers some of the technical challenges of collecting serial-based traffic. It also provides some advice on "quick-wins" that can be implemented in the shorter term as well as some suggestions for further research on the topic. Danagoulian pointed out that "you can hack electronics, but you can't hack physics," and the serial communications found in control systems is one layer closer to that un-hackable realm of physics (MIT, 2018).

1.1. Coming to Terms

This paper avoids the use of the generic term "Industrial Control System" in favor of Critical Infrastructure Control System (CICS), which is more specific to the types of control systems considered in this research (control systems supporting critical infrastructure). General ICSs do not exist in the same threat environment as CICSs, and because of this, they will have drastically different risk profiles. Because of this difference in risk profiles cybersecurity professionals tasked with defending a plant that produces widgets may not find this research as practical as those defending a plant that produces nuclear power, as the costs of implementing such a system in a widget factory are much more likely to outweigh the benefits.

In the context of this paper, the term "serial traffic" describes the non-routable, point-to-point traffic that comprises a relatively large portion of CICS communication when compared to traditional IT systems. An example is a PLC controller communicating with remote I/O using a Fieldbus protocol over an RS-485 link. Routable traffic refers to the typical TCP/IP style Ethernet-based traffic found in CICS and

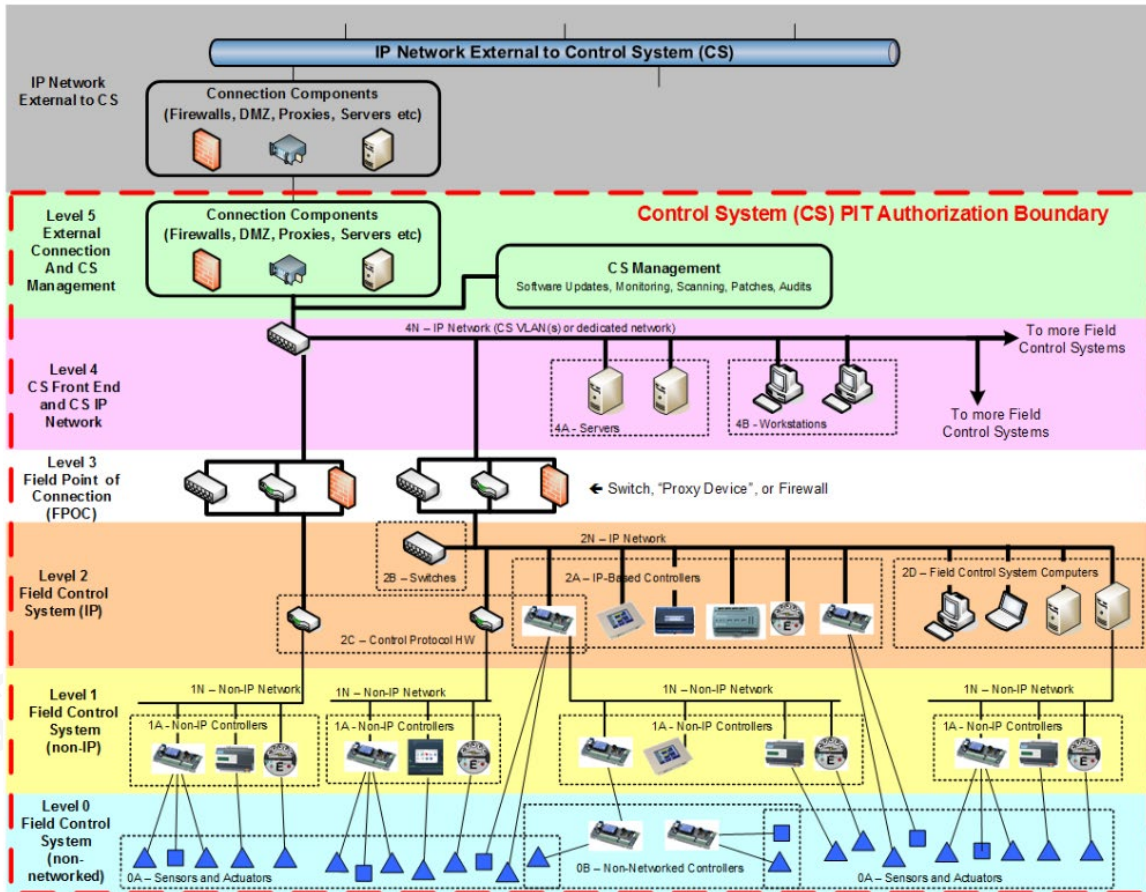


Figure 1: DoD 5-Level Control System Architecture

traditional IT systems. Examples of routable traffic include traffic such as TCP/22 SSH traffic, TCP/80 HTTP traffic, and even TCP/502 Modbus TCP. Figure 1 details the DoD's 5-Level Control System Architecture Model, which is similar to Purdue Reference Architecture, but bases the levels on communication types in use (Dalton, Gott, Oshiba, & McAndrew).

For instance, Level 0 is considered non-networked and contains non-networked devices that communicate with Level 1 and occasionally other Level 0 devices using digital and analog "hardware I/O" signals. Level 1 is defined as networked devices that do not use IP. Examples of this include devices using Modbus over RS-485. Level 2 devices are considered IP Networked devices, though they still use the Level 1N non-IP networks to communicate with Level 1 devices. The serial traffic discussed herein is within Level 1 and occasionally between Level 1 and Level 2 of the 5-Level Control System Architecture. Level 1 and lower communication traffic distinguishes control systems from typical IT systems; it bridges the gap between the cyber and physical worlds and is a defining reason why CICS systems require a different approach from the IT-system traffic collection products commonly offered by security vendors. The traffic within levels three through five is more akin to traditional IT traffic and is collected and analyzed in a similar fashion to IT systems.

1.2. Current State of Traffic Monitoring in Critical Infrastructure Control Systems

Many CICS organizations are only collecting routable network communications for analysis, such as that found in Level 2 and above. In contrast, non-routable system communications can make changes to the system and report system status without any analysis engines or analysts seeing the traffic involved (Dalton, Gott, Oshiba, & McAndrew). This is notable because while the communication traffic generated in a control system environment contains a much higher portion of serial-based traffic than that found in a typical IT installation, traffic collection is too often treated the same as it would be in an IT installation. Though some specialized vendors like Cynalytica and SEL are beginning to develop products that collect and analyze serial-based traffic, none of the top network visibility vendors offer turn-key solutions to address this traffic in their primary product lines. There is also a lack of rigorous discussion in the industry dedicated to making a risk-based decision on whether or not an organization should collect the serial-based traffic of their CICS for analysis. This seems to stem from the fact that the available tooling is specific to routable traffic and engineers' confidence that

malicious traffic will be found within the routable traffic before serial traffic becomes significant. This paper provides a brief case study of Stuxnet, a well-documented CICS attack, focusing on providing decision-makers with information that is based on actual scenarios that will assist them in making informed risk-based decisions regarding supplementing the collection of routable traffic with serial-based traffic. The feasibility of capturing this traffic and the expected benefits of collecting it is investigated to this end. Some other notable CICS attacks are mentioned in order to help more clearly define some arguments, but they are not discussed at length as a case study.

1.3. Shortcomings in the Current State/Defining the Situational Awareness Gap

There are three major contributors to the process-level gap in defenders' situational awareness within a CICS: (1) defenders rarely have the knowledge necessary to effectively monitor and defend at the process level, (2) high-quality tools are challenging to build due to the highly specialized and often proprietary nature of communications at the process level, and (3) the traffic between Level 0 and 1 devices is rarely, if ever, collected. Even if organizations began collecting serial traffic tomorrow, the problem would not be solved. Closing the situational awareness gap requires collaboration between the process engineers and cybersecurity professionals that often find themselves at odds. This will ensure that decisions made regarding the process are fully informed. Collaboration between security professionals, device vendors, and security tool vendors is also essential. Without this collaboration, the tools that give visibility into the proprietary protocols used in these control systems will be overly challenging to build and will potentially violate licensing agreements. This collaboration will allow for effective designs of process monitoring architectures that provide visibility into the legacy protocols. These protocols tell a more accurate story about what plant equipment is being told to do and what it is actually doing.

In the article, "*To Kill a Centrifuge*," Ralph Langner rightly points out that CICS targets in the US will likely not be highly specialized government facilities but typical critical infrastructure systems that are configured and operated in consistent ways with

standardized equipment (Langner, 2013a). It is easier for adversaries to perform intelligence gathering on these systems and they are well understood by many control systems professionals. This is more of a liability to the security of CICSs than the encroachment of Ethernet-based networks toward Level 0 and is why organizations should consider the serial gap in situational awareness as part of any robust and mature CICS security program's risk calculations.

1.4. Communications and Protocols Contributing to the Situational Awareness Gap

COM (communication) ports, as shown in Figure 2, are a common mode of serial traffic communication in CICSs (Hope, 2020). These are typically used for communications between controllers and workstations, between two or more different controllers, and between controllers and remote I/O.




Figure 2: DB9 Serial COM Port

The protocols used in serial communication are sometimes open protocols, like PROFIBUS, however, it is common to find proprietary legacy protocols in use or even open protocols that carry proprietary protocols as a payload. Even if this traffic can be collected, parsing it for analysis is not a simple task; custom parsers must be written, and the intricacies of the expected traffic must be well understood. Because of the lack of commercial off the shelf (COTS) solutions for collecting, parsing, and analyzing this proprietary traffic, it is often not collected or analyzed.

An often-overlooked form of serial communication within CICSs is the console




Figure 3: *Siemens S7-417*

port of devices used to communicate with the management plane of networking gear throughout the plant. This is an example of serial traffic that can be found at all layers of the Purdue reference architecture. Traffic over this port is rarely, if ever, collected and analyzed.

An example of these protocols and connectors in the field is the Cascade Protection Systems (CPS) at the Natanz nuclear facility in Iran. The CPS used PROFIBUS to communicate with Siemens S7-417 controllers. Figure 3 shows the DB9 port that can be used to communicate with the CPS over PROFIBUS. This example will be further discussed as part of the case study in

Section 2 because the Natanz CPS was targeted in the Stuxnet overpressure attack on the uranium enrichment centrifuges.

2. Research Methods and Event Analysis

The presentation of this research is targeted to IT and OT cybersecurity engineers, process engineers, and project managers in the critical infrastructure control system arena. This excludes control systems that would not be considered "critical

infrastructure" based on the Presidential Policy Directive on Critical Infrastructure Security and Resilience (PPD-21) because these systems will have vastly different risk considerations (White House, 2013). Since the attacks analyzed in this report are attacks against the energy sector, this analysis should be considered biased toward the energy sector.

Several events were considered for this research, including Stuxnet, Triton/TRISIS/HatMan, HAVEX, BLACKENERGY2, and CRASHOVERRIDE since they were targeted toward control systems, and were at least partially documented. Ultimately, Stuxnet was selected for this research for several reasons. The attacks within the Stuxnet campaign have been copiously written about, so there is a wealth of unclassified information available. Stuxnet is also well known outside of the critical infrastructure community and therefore is more relatable to those outside of the cybersecurity field, one of the groups to which this research is targeted.

2.1. Stuxnet

The targeted campaign outcome of Stuxnet was to disrupt Iran's nuclear capabilities by crippling the uranium hexafluoride gas centrifuges located in Natanz, about 150 miles south of Tehran. Because Stuxnet was specifically engineered to affect the process without drawing the attention of plant engineers or security personnel the problems it created at Natanz would have been difficult to distinguish from typical operations error (Langner, 2013a). This inability to distinguish between typical errors and malicious activities is common across CICS attacks (Weiss, 2010), and it reinforces the importance of having visibility further into the process in order to determine what is actually happening within the process. The history of Stuxnet is avoided in favor of focusing on attack specifics that can be used to discuss the costs and benefits of collecting and analyzing serial data.

2.1.1. Two Attacks in One

The mode of the Stuxnet attack was changed during the course of the campaign, and each method will be treated separately in this discussion. The first attack targeted the overpressure system, which focused on an overengineered dump system required at Natanz for the centrifuge system to operate and overcome the inability to manufacture parts to the exacting requirements necessary for the equipment. While an extraordinary feat of engineering, this mode of attack was very fragile and is likely to stop working in the event of small changes being made to the process (Zetter, 2016).

Later in the campaign, the noisier but more robust drive speed attack was deployed. It targeted the variable frequency drives in the system controlling rotor speed and was used to physically stress, weaken, and occasionally destroy specific centrifuges within groups of centrifuges called cascades. It is necessary to use cascades in the enrichment of uranium because a single centrifuge is not capable of separating the uranium (product) from the waste, so they are grouped into stages that are piped together in a cascade. If the right centrifuges within the cascade are successfully targeted, the process must be aborted. The overpressure attack did not rely on IT systems as much as the drive speed attack, so these attacks will be treated separately in the following sections.

2.1.2. Overpressure Attack

The first attack manipulated isolation valves and pressure sensors controlled by S7-417 controllers (Figure 3) that were used to maintain pressure within centrifuges at a level below that which they were initially designed. This was necessary due to Iran's inability to manufacture the centrifuges at the level of precision required to produce the intricate parts that the centrifuges rely on. The attackers used this foothold to manipulate the physical process (Langner, 2013b).

Control and situational awareness were both subverted in this attack by placing malicious code between the process and the legitimate PLC logic. This malicious code recorded normal operations traffic for a period of 21 seconds and later replayed the

normal traffic to both the legitimate code within the controller (which would still be running) and also to the Human Machine Interfaces (HMI) screens. Though the manipulation of view to the HMI is ostensibly the big problem here, the alarm systems were also being fed fabricated data that indicated operation within acceptable parameters. Loss of situational awareness was also achieved through a de-calibration process to prevent over-pressure safety valves operated by dedicated pressure controllers (but connected to S7-417 controllers via a datalink) from opening when they should have. This de-calibration process would also result in the pressure controllers showing normal values on their display panels when abnormal conditions were present (Langner, 2020).

The PR 4000 controller that was used to manage pressure at Natanz takes an electrical signal and correlates it to a specific pressure, with different electrical signals indicating different pressures. The mapping between these electrical signals and actual pressures must be calibrated, so that a sensor sending a specific electrical signal to indicate a pressure is mapped to the correct pressure by the controller. As part of the overpressure attack, this mapping was manipulated so that any electrical signal from the sensor was mapped to a normal value, even if the pressure raised beyond expected operating values.

The Cascade Protection System (CPS) was the general target of the Stuxnet overpressure attack, and like other critical infrastructure control systems, it used vibration monitoring to detect components operating outside of their designed parameters. The ineffectiveness of this protection system at detecting attacks indicates that relying on these protection systems alone to detect anomalies is insufficient and the added visibility into the process from serial collection may provide valuable data to increase situational awareness.

2.1.3. Drive Speed Attack

The drive speed attack did not use HMI manipulation to trick operators, the overpressure attack was the only one of the known Stuxnet attacks to use the recording and replaying of typical operations data. Legitimate code was merely suspended when the

attack code was in use. The operator displays are fed from controller memory rather than by information from the field devices themselves. This memory must be actively updated by the logic, which is interfacing with the field devices (Langner, 2013a). If the logic that typically does this is suspended during an attack, the expected steady-state value of the drive speeds would be seen rather than the actual values that would be seen changing during the attack.

Consider the frequency converter architecture likely used at Natanz and shown in Figure 4 (Zetter, 2011). This architecture can be compared with that of Figure 8 in Section 4. It details a PROFIBUS DP architecture which the Siemens CP-342-5 is commonly used to implement. Although the operators would see a static value for drive speed, collecting traffic destined for the CP-342-5 communications process would have revealed the fluctuating drive speeds. Though the changing drive speeds would have been audible, the discrepancy between what the frequency drives were reporting and what the controller was reporting would have pointed analysts to the controller.

WinCC is software developed by Siemens that is used to communicate with, and gain vision into, the ICS process. Researchers believe that the drive speed attack was geared toward a system where WinCC software was in use, but they have yet to uncover any hard evidence that WinCC was in use at Natanz. Langner has shown that if Stuxnet had used WinCC to interface with the controller, the unmistakable TCP/IP attack traffic would be found on the network (Langner, 2013a). If WinCC was in use and generating this well-known TCP/IP traffic, then this situation is an example where detecting the drive speed anomalies and linking them to malicious intent could have happened at the IT

level. Therefore, if that information was available at the IT level, then the engineers at Natanz would have missed the traffic simply because they were not looking.

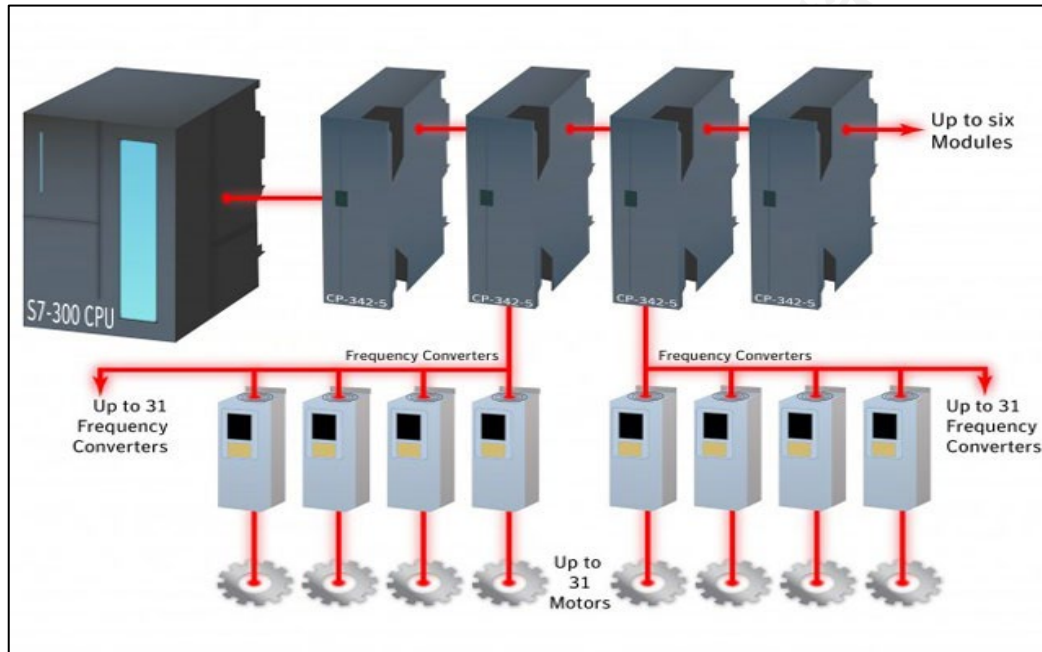


Figure 4: Suspected Natanz Frequency Converter Architecture

3. Findings

3.1. CICSs Require Specialized Considerations

As Langner notes in the executive summary of his seminal work "*To Kill a Centrifuge*"; to suggest mitigation strategies such as air gaps or robust anti-virus and patching programs implies a misunderstanding of the lessons learned from such a cyber-physical attack (Langner, 2013a). The solution to this problem is not rooted in the routine; it requires a more specialized approach, which should include, among other things, the collection of serial-based traffic. The communications are less noisy and more predictable in the serial space— furthermore, collecting this data decreases the options for the attackers and increases the opportunities for the defenders, which is valuable on both accounts.

Langner lays out the three layers involved in a cyber-physical attack: (1) the IT layer, where nearly all of today's CICS cybersecurity is focused; (2) the control system layer where process control is manipulated; and (3) the physical layer where equipment, the environment, and human life are physically put in jeopardy. TCP/IP communications form the bulk of traffic in the IT layer, while serial communications play a larger role in the two lower layers, which distinguish control systems from standard IT systems. In a paper written in the early 1990s, Bill Cheswick described a problem facing organizations' network architecture at the time. It described a "hard shell" that protected the network from the outside, but if an adversary were to move beyond this outer shell, they found themselves in the unprotected "soft and chewy center" of the network (Cheswick, 1990). To ignore these the lower layers of communication within a CICS environment is suggestive of the "Soft and chewy center" described by Cheswick in this paper.

A common tactic, techniques and procedures (TTP) activity for a nation-state cyber-physical attack is the use of undisclosed vulnerabilities, which may bypass detection and protection measures placed in the IT layer. The use of these undisclosed vulnerabilities means that a defense strategy relying entirely on monitoring traffic activity at the IT level is inadequate when trying to achieve a high level of situational awareness. This strategy is inadequate because the actions associated with the malware may not become evident until they reach the control system layer, where operations begin to deviate from the expected. Another common TTP is the use of highly tailored procedures and tools, making signature-based detection less than ideal.

Loss of control, such as that seen in the drive speed attack, could also be detected by monitoring serial traffic at Layer 1. If the expected operator or PLC commands are not seen downstream at the collection device after they are enacted, this could be treated as an event. This may seem like "too little too late," but it would be one indication that can be used to distinguish natural equipment malfunction from malicious manipulation. Distinguishing the difference between normal malfunction and malicious manipulation is valuable because without enough information they often look the same. As seen in the

overpressure attack, making an attack look like a natural failure allows an attack to be continuous without raising suspicion.

3.2. Issues Introduced by Lack of Visibility and Trust

Stuxnet's overpressure attack contained two examples of "loss of visibility." One was achieved through the malicious calibration of sensors in a way that the pressure controller, shown in Figure 5, would indicate normal pressures on its display and in conversations between itself and the control system, even as actual system pressures increased. Although after this malicious calibration was completed, serial traffic carrying pressure data between the S7-417 and the pressure controllers would appear normal, even if it were abnormal, the calibration process itself would have been apparent within the serial-based traffic. The second loss of view situation was brought about by the recording of normal operations traffic and replaying it to the legitimate code, which was still running throughout the attack. This blinded the system from situations that may have triggered alarms or may have even caused the safety system to step in and halt operations. As noted in Section 2, information made available through serial collection could have detected this situation.



Figure 5: PR 4000 Controller

Some other items that can occur as a result of a loss of view situation include misrepresenting the status of equipment; for example, a valve looks like it's closed, but it is open. The consequences resulting from such a situation can be found by examining the Three Mile Island event that occurred on March 28th, 1979. One of the lessons learned from this event focused on a valve in the primary system that was designed to

automatically open when exposed to high pressure and close again once the pressure was adequately relieved. At TMI-2, this valve was stuck open, but the operator display did not receive signal showing that it was in the open state. This led to cooling water escaping through the stuck valve and ultimately to the reactor overheating. This illustrates the real-world effect that a malicious actor could theoretically cause by manipulating view in such a way as to trick the operator into doing something catastrophic or cause the Safety Instrumented System (SIS) to stand by when it should be taking action. The malicious actor could also cause the process to be automatically stopped by tricking the SIS into thinking that the process has gone so far out of bounds that it must be automatically stopped, even though it operating as expected.

3.3. Operational Uncertainty

Operational uncertainty can be just as damaging as an actual attack; if an operator is unsure of the process state, then they may need to halt the process out of an abundance of caution. As noted in Dragos' TRISIS analysis, the manipulation of the SIS can create operational uncertainty, which could lead to a reduction or stoppage of operation (Dragos, 2017). The increased situational awareness resulting from the collection and analysis of serial data in Layer 1 would not only provide earlier detection, but it would also help to speed investigation and troubleshooting of unexpected values, thereby allowing a return to normal operations in a more expedient and safe manner. In the event that the SIS causes the process to be halted completely, the collection of serial traffic can help to illuminate the cause of the safety system's intervention and halting of the process, allowing for a more confident return to normal operations.

3.4. Field Device Manipulation

Serial collection to detect direct field device manipulation at Level 0 is less likely to have as extensive of an impact as the collection at Level 1. Level 0 communications consist mostly of analog signals such as 4-20 mA, 0-10 V or simple digital on/off signals that are not transmitted using serial-based protocols. An example of this is a pump telling an analog to digital converter that it is open at 50 percent by responding to a 20 mA

signal with an eight mA signal. Because these signals are basic digital and analog signals there is nowhere in the communication path that would gain value from the collection of serial traffic.

3.5. Bypassing Level 3, 4, and 5 Controls to Manipulate Fieldbus Traffic

A brief description of some ways in which attackers can bypass controls based on the TCP/IP Ethernet world follows. The below examples demonstrate situations in which the collection of serial traffic would be valuable because the collection would provide visibility into events that would not be easily seen by Ethernet-based collection systems.

Supply chains can be compromised. The Stuxnet attack showed evidence of this at a frightening level with the compromise of valid digital signing certificates (Zetter, 2016). However, the compromise of supply chains can also affect hardware. A nation-state that produces PLCs could compel a company to place hardware backdoors or undocumented malicious code into the PLCs logic.

Similar to clandestinely installed code by a manufacturer, a rootkit, by design, would exist on a system without the knowledge of system maintainers or operators. From this vantage point, a rootkit could manipulate the process without tripping IT-based sensors.

A malicious insider with physical access could load malware onto a system by simply plugging a device into a controller that can overwrite the controller's logic. This would be even easier if there were an available console port. However, if the insider were trusted, it would be trivial for them to unplug a cable and plug in the malicious device for a long enough period to overwrite the logic and then replace the original cable.

A common thread of attacks on critical infrastructure control systems has been the use of malware that takes advantage of unknown, or zero-day vulnerabilities. By definition, the world has had no exposure to this malware, so existing signatures offer no immunity. Behavioral-based detection in a deterministic environment, however, is very

effective. The more representative the data set is of the entire system, the more effective behavioral based detection is, and so the collection of serial data would increase the value of behavioral-based detection (Falliere, Murchu, & Chien, 2010).

Vendors often require direct access to their products that are deployed in the field. Even with DMZs, these products are accessible by the vendor, so if a vendor is compromised by a sufficiently clever attacker, they can access the vendor equipment directly over a channel that expects Ethernet-based traffic to be accessing the system. Compromise of a vendor could also lead to malicious firmware updates that include rogue code. The most recent example of this can be seen in the Solarwinds supply chain hack.

3.6. Making the Risk-Based Decision

The collection of serial traffic is shown to be effective and has merit, however, it is not a replacement for getting the common critical controls and basic security hygiene right first, like the Level 1 configurations described in the CIS Benchmarks for the technologies that are deployed in the plant. CIS Benchmarks are available for most major operating systems and software applications. Figure 6 shows the ARC ICS security maturity model, which was designed by the ARC Advisory Group, a group formed in the 1980s that focuses on operational technology. An organization should at least be in the "Manage" portion of this maturity model before considering the collection of serial traffic, but serial collection would provide the most significant benefits to organizations in the "Anticipate" stage, where anomaly and breach detection are already ingrained in the organization's processes. The application of other security controls should not be compromised in order to implement serial collection and analysis. Existing

controls should be fully and tightly integrated before considering the addition of serial-based collection.

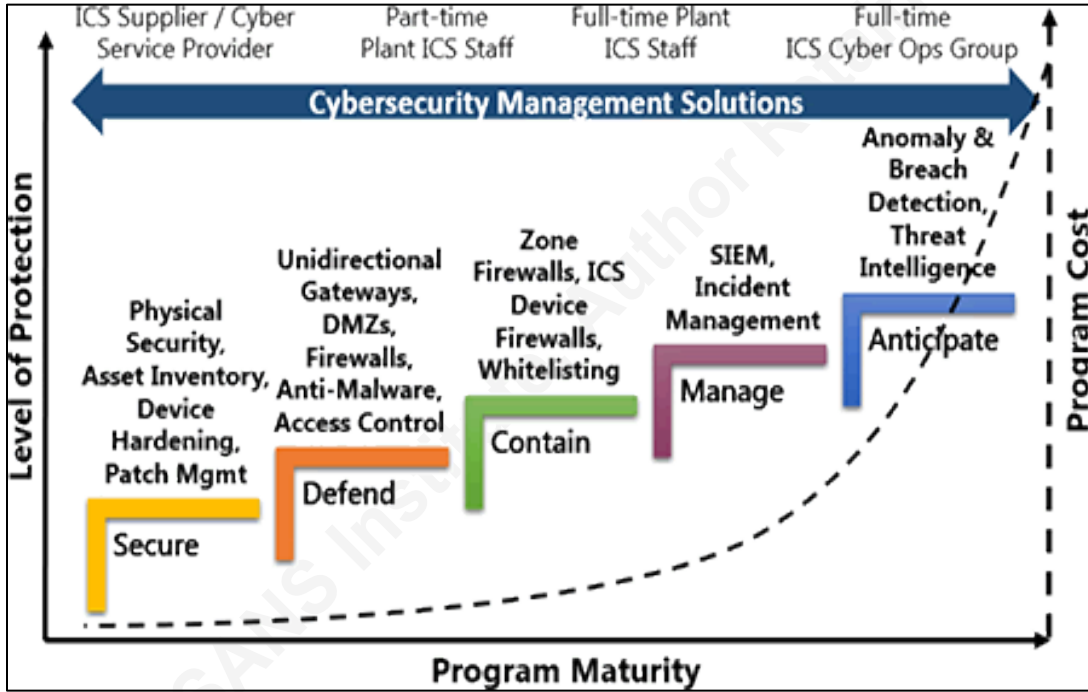


Figure 6: ARC ICS Security Model

As control systems trend toward Ethernet-based communications in Levels 1 and 2 of the DoD's 5-Level Control System Architecture, the effectiveness of collecting serial traffic to improve situational awareness will decrease. This decrease is expected because this communication will then be picked up by the ubiquitous TCP/IP-based collection and analysis systems. Organizations must look to the future of their industries, however slow that future may develop. The AP1000 reactor design by Westinghouse Electric Company is the latest reactor to be put into commission generating power (Sanmen, China). Even the latest technology used in this AP1000 Nuclear Power Plant design relies on serial-based traffic (the first of its kind plant only began producing power in June of 2018). Considering the expected operational lifetimes of these types of systems, the Energy Information Administration puts the average age of nuclear reactors in the US at 40 years old (Office of Nuclear Energy, 2020), it is not too late to start thinking about the move toward the capability to collect and analyze serial-based traffic.

When communicating with management and other stakeholders, it is important to avoid overstating the dangers of not collecting serial data. While it is very beneficial to collect serial traffic, deciding against its collection in an organization with a mature security program is not necessarily an open invitation to disaster. When considering the risks of compromise and catastrophic destruction of equipment, the presence of items like mechanically based safety features must be considered. This includes items such as mechanical/spring actuated pressure valves that will relieve pressure prior to a catastrophic failure or passive cooling that is designed into a reactor vessel so that even with the loss of cooling pumps, the reactor core will continue to be cooled by natural physical processes and avoid the accidental melting of the core.

Finally, when discussing the possibility of serial traffic collection, it is important to bring the conversation beyond the benefits as seen from a cyber security perspective and frame it around operational wins as well.

4. Quick Wins

Both Windows and Linux operating systems have options for interfacing and monitoring Universal Asynchronous Receiver-Transmitter (UART) devices. UART devices are known as COM ports on Windows machines and TTYSX devices on Linux-based machines. Some COTS software exists that can collect and parse this data, and specialized software can always be written. Though it would not be real-time, this could be a way to collect information on serial communications without the addition of any new hardware. This data may be best treated as log data as it is formatted as character data.

In the case of the Natanz configuration, the S7-414 controllers communicated using PROFIBUS, but the PR-4000 pressure controller did not. This required a communication gateway that converted the specialized PROFIBUS communication to RS-232 (Langner, 2013a). This traffic could be picked up on the RS-232 side of the gateway and combined with the standard RS COM traffic analysis to expand on the COM port communication "easy win." This particular architecture can be seen in Figure 7.

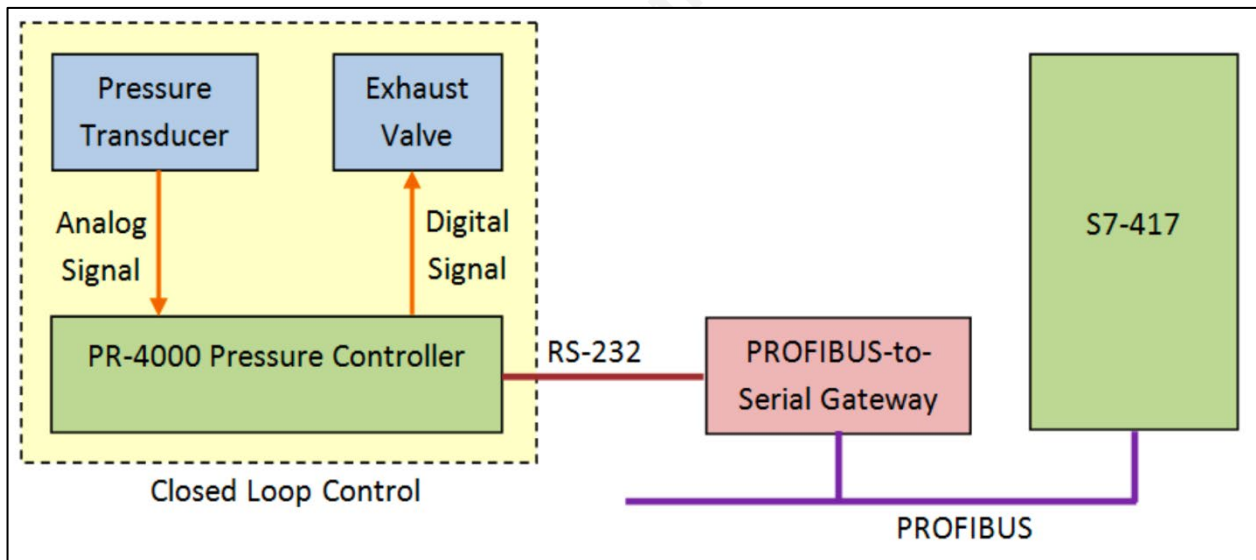


Figure 7: S7-417 to PR-4000 Pressure Controller Interface

Profibus DP is a deployment of PROFIBUS that decentralizes the sensor and actuator I/O from the PLC. This allows the I/O to be installed closer to the field devices, simplifying cabling and relieving space constraints within controller cabinets. Decentralization is achieved by installing I/O modules close to the process, separate from the PLC, but using the same field cabling as if they were installed with the PLC. The same collection of I/O modules as a standard PROFIBUS deployment is used; however, an extra module is added that is used to send combined data from the I/O modules over an RS-485 link to the PLC. A link like this typically uses RS-485 rather than RS-232 because of the distance capabilities of RS-485, which supports runs of 1200 meters vs. 15 meters of RS-232 (Frenzel, 2013). This is an excellent example of a chokepoint that can be used to collect serial data from many devices at a single location (Langner, 2013b). A

possible collection point at this location could be achieved using a COTS PROFIBUS "PG" style connector that allows for a PROFIBUS cable to be connected to the PLC and provide a DB9 port for another cable to piggy-back onto. Figure 8 shows the PROFIBUS DP architecture and the PG "piggy-back" connector. A benefit of this approach is that the controller and interface were designed with this type of configuration in mind. The field devices on the bottom of the figure connect to the remote I/O in the center of the image. The remote I/O connects to a DB9 port on the PLC in the top left via an RS-485 serial connection using an interface module. This DB9 port (red arrow) is where the PG connector would be connected. Collecting serial traffic at this location would require attackers to manipulate data at the remote I/O collection points or at the field devices themselves in order to adversely affect the situational awareness into the state of the process.

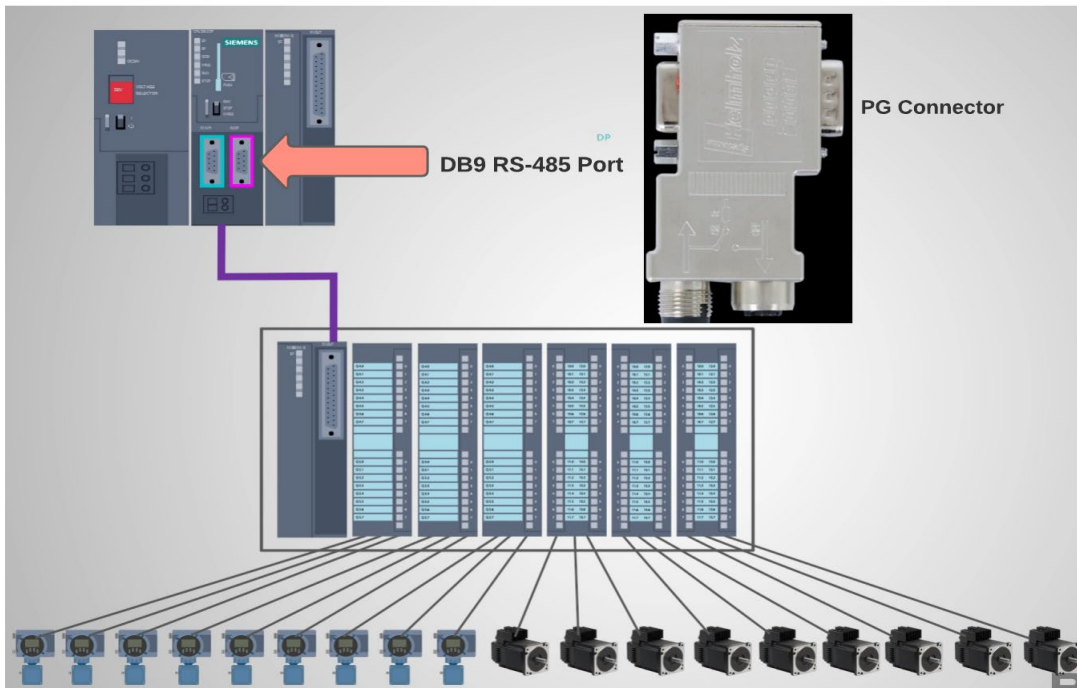


Figure 8: PROFIBUS DP and PG Connector

5. Future Research

While industrial control systems have been being engineered for many years, the concepts and ideas behind securing them are relatively new. Because of this there are plenty of unexplored opportunities to pursue research based on the preliminary examination posited here.

Both Stuxnet attacks relied on creating a situation where the operator and engineers had lost situational awareness. Fitting sensors that can monitor process status based on expected physical signatures (temperature, vibration, or humidity, for example) may be worth further research. An example of this would be an automated fluid valve that can be set to several positions. Measuring the vibration of the valve when fully open and again when it is fully closed would almost certainly create different vibration "fingerprints" or profiles. If these known vibration profiles were used to compare against vibration profiles coming from the field, it could be used to determine the integrity of the sensor data. If the valve is reporting that it is closed, but the vibration signature coming from the valve sensor looks like it is set at 50% aperture, there is a chance that the process is in a loss of view state, and it should be investigated. Though this too can be overcome by an attacker, it is a significant step in moving past the "soft and chewy center" that now exists in CICSs by adding another layer of security to the architecture.

The concept of comparing expected physical signatures can also be used to more accurately predict equipment failure. In a paper for The Institute of Electrical and Electronics Engineers (IEEE), Chen et al. describe a relay-assisted wireless sensor network that uses Kalman filtering to estimate unknown physical parameters (Chen et al., 2015). It is possible that sensors deployed in this manner could be installed without interruption of the process, which would be a boon to legacy systems that are difficult and costly to shut down. The results of the Kalman filter can be reconciled with the process status reported by the control system to detect possible loss of situational awareness and component degradation through consensus-based voting, similar to the tactic used in safety systems.

The creation of tools and rules related to the collection and analysis of serial-based traffic would be a great help to the community. Aron provides a process for creating custom Snort rules for OT environments that engages OT engineers (Aron, 2020). This process may be able to be adapted to analyze and generate alerts based on the highly custom protocols often found in CICSs. These tools could consider how to normalize the serial data so that analysis can be more efficient.

The following list provides more opportunities for future research to prove viability in a lab setting.

- Using collection points like PROFIBUS PG connectors to collect data from the network that can be analyzed
- Collecting serial traffic from UART ports on Linux or Windows workstations

In forensic science, Locard's Exchange Principle states that any interaction with a crime scene will leave some kind of trace of that interaction, even a non-malicious interaction. This holds true for computer systems as well. Windows operating systems are the most common systems in use within ICSs, and communications between these Windows systems and Level 1 devices occasionally occur over serial links. Considering Locard's Exchange Principle, studying how this communication interacts with native Windows artifacts can identify the ways in which serial data changes artifacts like Windows event logs or the registry. This can be used to gain more insight into what is happening on serial ports.

6. Recommendations and Conclusions

As noted by Langner, although the Stuxnet attack was an attack that was highly focused on a single mission, the TTPs were reusable. When making risk-based decisions regarding the types of traffic to collect, analysts should update their threat model with common TTPs for cyber-physical attacks that are significant to their particular

organization to help determine if the collection and analysis of serial-based traffic would be beneficial.

If implementing serial-based traffic collectors, be sure to isolate them from the serial bus using data diodes or similar gap technologies, such as Cynalytica's SerialGuard. This will prevent collectors from becoming a point of access into the low-level ... and eliminate the risk of collectors injecting unwanted or disruptive traffic into the control system network.

The specialized nature of the process and protocols at the level where collection and analysis of serial information would be worthwhile requires the collaboration of engineers who are familiar with the process and with security controls, processes, and architectures. Cultivation of the relationships between these teams should begin about the same time that an organization starts to consider the collection of serial traffic. These relationships can take a long time to develop, and even if serial traffic collection is not pursued, these relationships will prove invaluable.

Due to the fact that many legacy protocols are used in Levels 1 and 2 of control systems, vendors need to be heavily involved in developing parsers to analyze the serial traffic. Organizations should pursue relationships with vendors and advocate for more access into these legacy protocols. Also, system architects should look for vendors that are already integrating with systems in use at their organization's facility. For instance, Emerson Automation Solutions and Dragos have recently teamed up to integrate Ovation, a popular distributed control system often found in water systems and power plants, into the Dragos platform (Emerson, 2019)

Relying exclusively on existing monitoring systems such as vibration monitoring is not sufficient because these systems can be bypassed, and without looking at serial data, it is harder to tell if the information being reported by these systems is accurate.

Avoid limiting consideration of serial traffic value to conditions existing during an attack. Items leading up to an attack, such as the de-calibration of the PR-4000 controllers, are arguably more important to consider.

A robust cybersecurity program for a CICS must go beyond merely making a list of assets, tracking them, and applying IT-based security controls to secure them. A robust program should define physical vulnerabilities and analyze the different points where monitoring can be employed to view the process accurately.

Some of the lessons learned from Stuxnet were applied to non-specific control system configurations; however, the thought processes and techniques laid out can be applied to and considered for an individual organization's installation and threat profile, allowing them to make more data-driven decisions regarding the collection of serial network traffic in their critical infrastructure control system.

References

- Aron, A. (2020). 60870-5-104 protocol snort rule customization (Master's thesis, SANS Technical Institute, 2020). Bethesda: SANS Institute.
- Chen, C., Yan, J., Lu, N., Wang, Y., Yang, X., & Guan, X. (2015). Ubiquitous monitoring for industrial cyber-physical systems over relay-assisted wireless sensor networks. *IEEE transactions on emerging topics in computing*, 3(3), 352-362. doi:10.1109/tetc.2014.2386615
- Cheswick, B. (1990). *The design of a secure internet gateway* (AT&T Bell Labs, 1990). Murray Hill, NJ: AT&T Bell Laboratories.
- Dalton, J., Gott, J., Oshiba, E., & McAndrew, M. (n.d.). *Cybersecurity of facility-related control systems* (United States of America, Department of Defense, Assistant Secretary of Defense). Department of Defense.
- Frenzel, L. (2013, April 16). What's the difference between the RS-232 and RS-485 serial interfaces? Retrieved December 01, 2020, from <https://www.electronicdesign.com/technologies/communications/article/21800966/whats-the-difference-between-the-rs232-and-rs485-serial-interfaces>
- Falliere, N., Murchu, L., & Chien, E. (2010). *W32.Stuxnet dossier* (Symantec, 2010).
- Greenberg, A. (2020, October 23). How 30 lines of code blew Up a 27-Ton Generator. Retrieved December 01, 2020, from <https://www.wired.com/story/how-30-lines-of-code-blew-up-27-ton-generator/>

- Hope, C. (2020, June 30). What is a Serial Port? Retrieved December 01, 2020, from <https://www.computerhope.com/jargon/s/seriport.htm>
- Langner, R. (2013, November). To kill a centrifuge [Web log post]. Retrieved from <https://www.langner.com/to-kill-a-centrifuge/>
- Langner, R. (2013, November 19). Stuxnet's secret twin. Retrieved December 01, 2020, from <https://foreignpolicy.com/2013/11/19/stuxnets-secret-twin/>
- Langer (2020, July 23). The Stuxnet story: Data-driven OT/ICS security. Retrieved December 01, 2020, from <https://www.langner.com/2020/07/the-stuxnet-story/>
- Shapournia, S. (2020, January 13). What is Profibus-PA and How Does it Differ from Profibus-DP? Retrieved December 01, 2020, from <https://realpars.com/profibus/>
- Weiss, J. (2010). *Protecting industrial control systems from electronic threats*. New York: Momentum Press.
- Zetter, K. (2011, July 11). How digital detectives deciphered Stuxnet, the most menacing malware in history. Retrieved December 01, 2020, from <https://arstechnica.com/tech-policy/2011/07/how-digital-detectives-deciphered-stuxnet-the-most-menacing-malware-in-history/5/>
- Zetter, K. (2016). *Countdown to zero day: Stuxnet and the launch of the world's first digital weapon*. New York: Crown.
- Dragos. (2017, December). TRISIS malware: Analysis of safety system targeted malware.

Emerson. (2019, July 29). Emerson selects Dragos to collaborate on cybersecurity protection for power and water industries. Retrieved December 01, 2020, from <https://www.emerson.com/en-us/news/automation/1907-dragos>

Massachusetts Institute of Technology. (2018, April 19). For nuclear weapons reduction, a way to verify without revealing: New isotope-detection method could prove compliance but avoid divulging secrets. ScienceDaily. Retrieved December 1, 2020 from www.sciencedaily.com/releases/2018/04/180419130910.htm

Office of Nuclear Energy. (2020, April 16). What's the lifespan for a nuclear reactor? Much longer than you might think. Retrieved January 15, 2021, from <https://www.energy.gov/ne/articles/whats-lifespan-nuclear-reactor-much-longer-you-might-think>

The White House, Office of the Press Secretary. (2013, February 12). Presidential Policy Directive -- Critical infrastructure security and resilience [Press release]. Retrieved December 1, 2020, from <https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil>