



SANS Institute

Information Security Reading Room

Industrial Traffic Collection: Understanding the implications of Deploying visibility without impacting production

Daniel Behrens

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.

Industrial Traffic Collection: Understanding the implications of Deploying visibility without impacting production

GIAC (GICSP) Gold Certification

Author: Daniel Behrens - danbehrens@gmail.com

Advisor: *David Hoelzer*

Accepted: May 27, 2020

Abstract

Due to the critical nature of industrial environments and the lifetime of deployed assets, many organizations do not have complete knowledge of what assets are operating in the environment and what communications are involved. With the continuous move to IP based communications for controls equipment, Cybersecurity continues to increase in importance and is a priority for many executives. Industrial controls are unique because they are interfacing with the real world, which has implications on human safety and the ability of an organization to maintain operations. Unfortunately, the criticality of these devices and the lack of robust network functions on many often requires the use of passive solutions to gather information. This paper will focus on outlining the potential impact of collecting network traffic, discussing the functions available on networking equipment to enable it, identifying possible deployment architectures and the pros and cons of each, and explaining a methodology to calculate the potential impacts.

1. Introduction

Industrial Automation and Control Systems (IACS) have an impact on many aspects of our everyday lives: the generation and delivery of electricity to our homes; the treatment of the water that we consume and use to clean; the traffic signals and crosswalks that keep vehicles and pedestrians safe; the manufacturing of the everyday items we use. All these examples and many others rely on the use of automation equipment to run, monitor, and maintain the processes involved. Historically, many of these devices leveraged siloed and proprietary communication methods that often leveraged serial communications. Standardizing on ethernet-based technologies has enabled more integrations, vendor options, and more interactions that have led to several enhancements and advancements in control and monitoring technologies. Unfortunately, the move to ethernet has led to an explosion of network-connected devices that are deployed for many years or even decades.

For many industrial assets, the addition of network connectivity was a mechanism to simplify deployment, to collect operational and status information, and to enable control of physically separated devices. However, these communication capabilities were often not developed with security features or robust networking functionalities as the focus. Many device vendors leverage open source or lower function IP stacks to achieve the desired functionality without significantly increasing their development complexity. Compounding the issue, many IP based networks in industrial environments have grown over time, exceeding the original design of the network and often without a focus on networking best practices. All of these factors have led to instances of attempts by organizations to scan or discover devices, causing device crashes or network failures that have stopped critical processes, impacted production, and created a concern or even fear of the impact of scanning or adding any additional traffic.

With the increase in connected devices, instances of malicious activity impacting production making the news, and the fragility of industrial assets and networks, a passive form of asset identification leveraging deep packet inspection and requiring full network captures has given rise in the industry. This research aims to outline the architectures, features, and functions available on network devices to provide full packet information, discuss possible impacts of these methods, and offer a simplified approach for organizations to calculate and identify which options are relevant. This research will not discuss the features, functions, or usefulness of any solutions that collect and analyze the traffic, as this is beyond the scope of this research and will remain neutral regarding any such solutions. The goal of this paper is also not to suggest or recommend any specific architecture; the goal is to convey the decision criteria an organization should consider when choosing options to capture network traffic.

1.1. Documentation Flow

The topics discussed in this document may be of interest to a wide range of users with varying backgrounds. Items such as industrial security and industrial network design can often cross the boundaries between information technology (IT) roles and operational technology (OT) roles. As a result, some discussions provide a basic overview of concepts to ensure that readers can make the most use of this research. Icons help the reader navigate the information. Readers can quickly pick and choose which topics are of interest or which the reader can skip due to already having a more in-depth understanding.



To depict topics that are IT-related concepts, such as network switching functions, the IT icon will appear next to the section heading.



To depict topics that are OT related concepts, such as "what is a PLC?", the OT icon will be located next to the section heading to depict this.

Sections that are relevant to both audiences will not have an icon.

2. Industrial Introduction

2.1. Industrial Devices

As previously discussed, industrial assets and automation devices control many vital processes that impact our everyday lives. One of the key characteristics that make an industrial automation and control device unique is the fact that it interacts with the real world. Some essential devices include:

- Programmable Logic Controller - PLC - Essentially, this is a small ruggedized computer typically based on real-time operation and used to receive inputs, leverage logic to make decisions, and trigger outputs.
 - Input/Output - I/O – This is a device that receives electrical signals (input) and sends electrical signals to trigger an action (output). A PLC commonly controls them.
 - Human Machine Interface - HMI – This is defined as hardware or software that provides an operator/user information regarding the status of the equipment or process and enables operators to interact with the machine / PLC.
 - Drive – This is a piece of equipment that interfaces with motors to create motion.

There are many different types of drives including variable frequency drives or VFDs, which can control the speed of a motor typically by changing the frequency of the electricity provided to the motor and servo drives, which control the exact location of the motor.

Different types of sensors can be used to gather information about the environment, equipment, material, or process for the use of control which includes but is not limited to:

- proximity sensor - used to determine if an item is present or the distance to an item
- temperature sensor - used to measure the temperature with varying accuracy depending on the type and material
- pressure sensor - used to measure the pressure or force applied to a device or material
- flow sensor - used to measure the rate of material passing by the sensor
- speed sensor - used to measure the speed of an item such as how fast a motor is spinning

Some other key terms that often come up in discussions regarding industrial controls are:

- Supervisory Control and Data Acquisition - SCADA - Typically a geographically dispersed system in which control is broken up into multiple sub-control systems
- Distributed Control Systems - DCS - control system often used in process applications where the control is distributed throughout a facility and is programmed to maintain and monitor set points.

2.2. Industrial environments

While there are many possible ways to define an industrial environment, for the purpose of this research, it can be defined as an environment where raw materials are combined or modified to create something new. These environments can range from

clean rooms where life-saving drugs are produced to coal electricity generation facilities, and can also range from environmentally controlled to dirty with substantial variances in temperature ranging from below freezing to above thousands of degrees Fahrenheit. These environmental factors can have a significant impact on the specific types of equipment needed and the need for communication between them.

Some simplified examples of industrial controls and environments include:

- A water treatment facility where flow sensors and level sensors measure the rate in which water and chemicals are entering large holding tanks and PLCs leverage those values to calculate how much to add of each before signaling to valves to close, which slows and eventually stops them from entering the tanks.
- A bottling plant where proximity sensors are leveraged to communicate at very high speeds to a PLC that controls a conveyor belt moving glass bottles until they reach a designated location. A filling machine quickly and precisely adds the right amount of product to the bottle, puts a lid on, adds a label, and moves along the line to be boxed and shipped out.
- An automotive plant where the frame of the vehicle moves from station to station as operators follow on-screen instructions on precisely which parts to install leveraging torque guns that automatically calibrate to the right torque based on which part is being installed.
- A steel facility where heat is generated by creating electrical arcs using high currents to melt down scrap materials before casting, forming, rolling, stretching, and treating based on the desired thickness and shape.

2.3. Industrial communications

There are many different protocols used across industrial environments depending on the type of equipment, function of the process, distributed nature of the control, and the time frame in which it was installed. For this research, we will focus on protocols which were either modified to run across standards-based Ethernet, or protocols that were developed for Ethernet from the beginning. While the industrial controls vendor can be one of the primary drivers for the protocol in use, process requirements can also dictate the type of protocol used.

Some examples of Industrial protocols that can be communicated across Ethernet include Ethernet Industrial Protocol (Ethernet/IP), Profinet, Modbus TCP, OPC UA, IEC 104, DNP3, BACnet/IP, IEC 61850, and a long list of vendor-specific protocols. Specific industries tend to use protocols that were developed with that industry in mind. For example, IEC104, IEC 61850, and DNP3 are examples of protocols often found in electrical utilities, whereas Ethernet/IP and Profinet tend to be more specific to manufacturing.

These protocols can be leveraged for reading I/O data by PLCs, collecting operational information such as production values for long-term storage and analysis, and HMIs to display and receive information from operators. Depending on the process, I/O communication could require responses at the speed of milliseconds or a critical piece of information that could be missed. These environments can involve flammable or even explosive materials where a catastrophic failure could lead to damage and loss of life. In many environments, operators are interacting with equipment, and failures could lead to injuries or worse. Engineers design and deploy solutions to safely fail and reduce the risk of these environments, but a loss of communications can still lead to damage and loss of product.

Daniel Behrens, danbehrens@gmail.com

2.4. Industrial Visibility

Due to the nature of the equipment and criticality of devices continuing to operate and communicate, any activity that could disrupt these is not going to be readily accepted and deployed. Thus, solutions that do not create any additional network traffic or attempt to communicate with critical assets are often preferred. These solutions often referred to as passive solutions, require the ability to collect all the network communications of devices on the network.

3. Networking Introduction

3.1. Network Traffic Collection

It is essential to collect the traffic as close to the end devices as possible, to have full access to the network traffic related to industrial devices. Many devices are only communicating to devices directly connected to the same switch, so attempting to collect traffic further up the architecture can leave gaps in what information is visible. Since this traffic is often used to attempt to identify resources on the network, full visibility is ideal for gathering and identifying as many devices as feasible. Several options will depend on the architecture and capability of the network devices that have been deployed to collect all the available traffic. For example, an unmanaged switch can be problematic, as an actual unmanaged switch with no management interface will typically not have a mechanism to copy traffic. This section will discuss the various options that are traditionally used to gain this visibility.

3.1.1. Network TAP

The network TAP is the only option that will be discussed in which it is not a function of a network switch. TAPs are devices that are deployed inline between the end device and the port of the switch the device is connecting. There are different types of TAPs with different levels of functionality. Ultimately, they all work by copying the traffic from the interfaces and sending it to an additional interface or interfaces that can be connected to the desired destination. The benefit of a TAP is that there is no dependency on the network equipment that is deployed, and typically there is little to no setup required. Still, the drawback is that each TAP is limited to the traffic to and from a single device, and thus, many TAPs will be required to gain full visibility. This always means installing additional hardware (the TAP itself), additional cabling, as well as possibly additional power depending on the type of TAP, selected.

3.1.2. Switch Port Analyzer (SPAN) or Port Mirroring

Port mirroring is the function in which the network switch itself is able to copy the network frames as they arrive to an interface or set of interfaces, and send it out another interface or set of interfaces without disturbing the original traffic. The methodology in which this is accomplished can vary from vendor to vendor, so it is always recommended to understand any possible limitations or impacts that are possible. Cisco refers to this technology as SPAN. The Cisco switches that were leveraged for this research perform this function in the switching hardware, which means there is no impact on the performance when port mirroring is leveraged. The advantage of leveraging port mirroring is that it can reduce the need to deploy a piece of hardware for every device to be monitored, and it can be enabled on the same switch in which the devices are already connecting. This also allows for the specific traffic that is being monitored to be more dynamic as it is a configuration change instead of a physical cabling change. In enterprise

switching, port mirroring is a reasonably ubiquitous function; however, in industrial switching, not all switches will have the ability to perform port mirroring. One downside to port mirroring is that it is locally significant, so organizations will need to either deploy a collection device for every switch or leverage devices capable of receiving mirrored traffic and sending it across the network. More information on the architecture can be found in section 3.2.

3.1.3. Remote Switch Port ANalyzer (RSPAN) or remote port mirroring

Much like SPAN, RSPAN is a Cisco-specific technology that is more generically referred to as Remote Port Mirroring. Much like port mirroring, remote port mirroring is a feature of the network switch in which the switch is duplicating the traffic of the configured source interfaces. The difference is that the switch is then putting the copied traffic into a specific Virtual Local Area Network or VLAN, which can then be sent across multiple switches. The benefit of this method is that it enables the collection of traffic from multiple switches and for it to be sent to a more centralized location, reducing the amount of additional hardware that may be required. However, one key consideration is the potential for this traffic to impact the other traffic on the network. The primary purpose of this research is to help organizations gain an understanding of the possibility of this additional traffic on the network, creating jitter and latency or competing for available bandwidth, and the amounts of traffic that cause the impact. While remote port mirroring is a feature on many enterprise switches, availability on industrial switches tends to be limited.

3.1.4. Encapsulated Remote Switch Port ANalyzer (RSPAN) or Encapsulated Remote Port Mirroring

While Cisco originally defined ERSPAN, unlike SPAN or RSPAN, ERSPAN was submitted as a standard and is a more commonly used term across the industry. Much like SPAN or RSPAN, ERSPAN is still a function of the switch in which traffic is copied from configured interfaces. The significant difference with ERSPAN is that a receiver for the traffic is defined, and the switch encapsulates the copied traffic into a Generic Routing Encapsulation or GRE header with the receiver as the destination. The benefit of this is that it allows for the duplicated traffic to be treated like all other traffic in that it can be routed with a destination IP address of the defined receiver. This allows for an even more centralized approach, but just like in the case of remote port mirroring or RSPAN, the traffic routing across the same network can cause jitter and could compete for available bandwidth.

3.2. Network Architectures

To capture all the traffic for the devices in the environment there are options that organizations will be able to leverage depending on existing infrastructure, physical layouts, and requirements of the solution that is receiving the captured traffic.

3.2.1. Distributed collection Devices




figure 1: distributed collection architecture

One option is to leverage TAPs or port mirroring on the switches in which the devices are connected. In this architecture, the number of collectors that need to be deployed depends on the specifics of the collection software and the physical layout of the environment. This option will require the most significant amount of new hardware to be installed and managed. Using only port mirroring or TAPs means any possible impact of the collection will only impact a single device or single switch.

3.2.2. Remote Port Mirroring or ERSPAN across the network




figure 2: Remote Port Mirroring / ERSPAN Architecture

With this option, the ability to send the captured traffic across the existing network is leveraged to reduce the number of collection points across the network. This deployment has the benefit of reducing the number of devices that need to be installed and managed. One consideration is the possible impact of the additional traffic. Sections 4.4 and 4.5 will discuss the methods an organization can look at to determine the possible amounts of traffic and determine the risk of impacting critical processes.

3.2.3. Out-of-band collection gear or network




figure 3: Out of Band architecture

In this option, any of the primary methods can be leveraged, but a network is deployed in parallel to the primary network to transport the captured traffic. This option has the benefit of allowing for a fully centralized data collection without the risk of the duplicated traffic impacting the control traffic. The downside to this approach is that it requires installing and maintaining a secondary network that will require additional cabling, power, and equipment.

4. Industrial Protocols

As discussed, there are many industrial protocols leveraged across industries.

For this research, we focused on two primary protocols: Ethernet/IP, and Profinet. While the nuances of different protocols will impact the values, the overall concepts and calculations will remain relevant.

4.1. Ethernet/IP

Ethernet/IP (Industrial Protocol) is an application layer protocol that works on standard TCP/IP (is the payload of an IP packet). Implicit and Explicit are the two primary communication methods used within Ethernet/IP.

Implicit messages leverage UDP unicast and multicast to communicate I/O information between PLC and I/O or between PLCs. This communication is typically local and involves persistent updates which are cyclical and based on a defined update rate. While the total packet size will vary based on the type of data and the number of data points, an implicit message is around 500 bytes long. The frequency of the message is application specific but can typically range between .5 milliseconds to 750 milliseconds with the default in Rockwell Automation deployments at 20 milliseconds.

Explicit messages leverage TCP traffic and are typically leveraged for interlocking or sharing data between PLCs, HMI to PLC connections, and administration of devices such as configuration changes and firmware updates. This communication will more commonly cross physical machines and possibly across a facility. Due to the nature of this communication, explicit messages may be ad-hoc or cyclical for polling the devices, but will often be 500 ms or higher. The size of an explicit message can vary based on the function and application but is typically around 1500 bytes.

4.2. PROFINET

Profinet has its roots in a serial based Fieldbus protocol called profiBUS but has extensions and features optimized for Ethernet-based communications. Three different communication channels are leveraged for data transfer between controllers and I/O devices.

PROFINET NRT or non-real-time is used for non-critical communications such as an HMI polling status information from a PLC. PROFINET NRT is built by leveraging TCP/IP communications and typically has a cycle time of greater than 100 milliseconds.

PROFINET RT or real-time leverages ethertype 0x8892 to bypass the IP stack and enable quicker communication. The result is cycle times as fast as ten milliseconds

which can be transferred over standards-based Ethernet. This is the primary communication channel leveraged by I/O traffic and is used in roughly 90% of Profinet deployments ("Five criteria for choosing an industrial ethernet - PI North America blog," 2009)

PROFINET IRT or Isochronous Real-Time leverages non-standard Ethernet features, including timing mechanisms, to provide sub-one millisecond cycle times. PROFINET IRT requires the network equipment to support it, which is not available on every ethernet switch.

4.3. Criticality of communications

As discussed above, in both Ethernet/IP and PROFINET, as well as the majority of industrial protocols, there is a concept of the control devices and PLCs requesting information from the I/O at a specific interval, relative to the application. The calculation of a timeout is specific to the protocol, the configuration of the controller and the device's vendor. Despite this, missing one of two cycles is often enough to cause an operation to react or even stop. As a result, that could mean as little as a 100-millisecond delay, or as little as a single packet drop could be enough to trigger a stoppage that requires human interaction to correct and can lead to loss of time and materials.

4.4. The impact of Jitter and Latency

Network latency is the time it takes for a request to get from the sender to the receiver. Jitter is the variance in latency across the network. While the goal is always to keep both numbers as close to 0 as possible, some factors can impact them. Overall network design, the distance of cabling, and the types of network devices leveraged can all impact the latency of the system. One of the goals of designing and deploying a system is to tune the application with these in mind. The time it takes for a request to reach an I/O and then the response to reach the PLC can be directly tied to how quickly

the machine can run and how short the timeouts of the application are configured. Unfortunately, this means that once the application is in operation, changes to the latency, or the increase in jitter, can make optimizing the equipment difficult and can cause operation stoppages.

When a network switch receives a frame, it reads the layer two information to determine where it needs to be switched. This small amount of processing is typically done using specialized hardware that can do it at the same rate as the information is being sent. However, if an interface is overloaded, the switch will begin to store or buffer the frames to avoid traffic loss. The more overloaded the interface, the more the traffic will buffer until either the buffer is filled and frames are dropped, or the traffic slows down and the interface can catch up. This process can cause an increase in the latency through the switch, and since it will most likely be variable, it will also cause an increase in jitter. This is where increasing the amount of traffic that is flowing through the network at any given time can have a real impact on the critical control traffic that is latency sensitive.

4.5. Overview of common amounts of traffic

The amount of traffic, which depends on the size of packets and the frequency or speed in which they are sent, is both protocol and application specific. As a result, it is impossible to provide a specific value that each organization can expect to see. This research aims to provide examples of common types of traffic and a set of calculations that organizations can leverage to approximate their environments.

5. Testing and Results

5.1. Overview of testing methodology

For this research, the testing was broken down into two scenarios. In the first scenario, a single switch is used, with the devices generating the traffic and the network capture all occurring on the same switch. A SPAN or port mirroring session was

configured to collect from multiple interfaces as the source and send the traffic to a single interface.




figure 4: Scenario 1 - Single switch architecture

In the second scenario, two switches are interconnecting with a 1 gigabit per second interface. On the first switch, a PLC is connected along with devices leveraged to generate traffic. On the second switch, an I/O block is connected (Which is being controlled by the PLC connected to the first switch), and then additional devices are added to generate traffic. An RSPAN or remote port mirroring session was configured on the first switch and sent on a VLAN that is carried over the interface between the two switches. On the second switch, a port mirroring session was configured to take traffic from the remote port mirroring VLAN and put out a single interface of the second switch. Lastly, the collection device is connected to the second switch.




figure 5: Scenario 2 - Multiple switch architecture

In both scenarios, testing was conducted by leveraging Ethernet/IP traffic and PROFINET traffic to see any differences in behavior between the protocols. The total number of transferred packets were observed and compared against the packets received by the collection device. For this testing, the collection device was simply a Raspberry PI running TCPDump.

To create the initial traffic capture and set a baseline for Ethernet/IP, a Rockwell CompactLogix L35E, and a 1734-AENTR Point I/O with an 8 point digital input and 8 point digital output was used. For the PROFINET baseline, a Siemens S7-1200 CPU and a Simatic ET 200 SP with an 8 point digital input and 8 point digital output cards was used.

To increase the amount of traffic, Raspberry PI devices running Ubuntu and leveraging SCAPY were used to manipulate the original traffic and replay packets onto the network. For each test, the amount and speed of traffic and the number of Raspberry PIs were increased until the connection between the PLC and I/O failed. To simplify the identification of when the connection is disrupted, the configuration for the I/O connection was to fault the processor.

5.2. Overview of results

In the first scenario, the only observed limitation is that the collection was unable to receive all the traffic when the aggregate traffic of the source interfaces exceeded the output speed of the destination interface. In other words, if leveraging a Gigabit Ethernet interface and pushing over a Gigabit per second worth of traffic, the collection interface would no longer see all the inserted traffic. However, at no point in time was the connection between the PLC and the I/O impacted by the collection traffic. The underlying architecture of a switching platform could provide different results, but the Cisco Industrial Ethernet switches that were leveraged for the test do not have any impact on the switched traffic when mirroring the traffic locally.

In the second scenario, in both protocols, there were two primary values for when interruptions could occur. First, when the amount of traffic captured on the first switch was large enough to cause congestion between the two switches. Second, when the actual traffic going switch to switch was increased large enough to cause congestion. It was found that the size and frequency of both Ethernet/IP and PROFINET traffic were

similar enough to each other and yielded similar results. Due to the fact that the 1 Gigabit per second interface between the switches is in fact full duplex, the switch to switch traffic's impact is per direction, whereas when capturing traffic and sending across via RSPAN, both directions are captured and sent across the link. As a result, the captured traffic can have a more significant impact much quicker. Ultimately the number of I/O per PLC will impact the amounts of traffic endpoints will generate.

5.3. Quality of Service (QoS)

Quality of Service or QoS is a mechanism that allows for the marking and prioritization of traffic as it leaves and enters a switch. While there are different versions based on the vendor and even within a vendor, it essentially allows for specific traffic to be prioritized over other traffic. When congestion occurs, the prioritized traffic will be dropped last, making it less likely to be impacted. In this testing, QoS was not configured and included in the experiment. One reason for this is that using QoS in industrial applications has had limited use. Often QoS is a more advanced function that not every vendor and platform supports and can add complexity to the switch and network configurations. Leveraging QoS would be a recommended function to be deployed in these environments, independent of any deployed monitoring solution.

Another reason for excluding QoS in this testing is that the goal was to identify the balance between gaining full visibility and not impacting production. Leveraging QoS would enable adding more devices without their traffic impacting the industrial communications but could start to lead to the dropping of the captured traffic and blind spots in the activity. To offer a clear set of values and derive accurate calculations, QoS was eliminated from testing.

6. Sizing Calculations

Based on testing, the average I/O connection in Ethernet and PROFINET is about 100 Bytes per packet with an average cycle time of 20ms. This means if we take 100 Bytes per packet at 50 packets per second, we will arrive at about .04 Mbps per I/O connection. A more accurate calculation would involve using the average number of I/O per PLC in the environment. If that data is not readily available, we can use the worst case, which is roughly 250 I/O per PLC. If we take 250 I/O connections, at .04 Mbps per I/O, that means each PLC is using roughly 10 Mbps.

To summarize, this puts the calculation at:

$((100 \text{ Bytes} * (1/\text{Cycle Time})) * \text{Number of I/O per PLC}) * \text{Number of PLCs} =$
bandwidth consumed by PLCs

Simply doubling this value provides an estimate of the amount of bandwidth that would be captured for all the PLC and I/O connected to a switch.

Using this value, an organization can calculate the amount of bandwidth performing a network capture will consume and compare this value to the existing bandwidth used on their switching uplinks. It would also be recommended for the amount of bandwidth consumed by network collection, and the bandwidth consumed by network traffic, to be less than 80% of the total capacity of a link.

7. Recommendations and Implications

7.1. Recommendations for Practice

Since network capture solutions can rapidly increase the total traffic crossing a network, it is highly recommended that while remote port mirroring or ERSPAN may be leveraged to reduce cabling and hardware deployment, it needs to be balanced to minimize the number of switches in which the collection traffic passes through. It is recommended that an organization take these numbers as an example, as more significant amounts of I/O, or more considerable amounts of data collection could impact the total amounts of traffic on the network.

7.2. Implications for Future Research

The testing performed was limited to I/O and HMI connections. With higher speed requirements such as motion or safety applications, the thresholds and impacts of traffic congestion could be even higher. With additional advancements such as Time Sensitive Networking (TSN), the drive to further tune and speed up process control continues to evolve.

8. Conclusion

While the scope of this research was limited to available equipment and hardware, the concepts outlined can help an organization determine which methodologies and architectures are best suited to fit their environments. As previously discussed, configuring advanced features such as Quality of Service, are strongly recommended to prevent congestion and jitter issues when there are large amounts of critical control traffic leveraging the network. It is imperative that organizations recognize the implications of monitoring industrial network traffic and understand the risks involved.

References

Sistrunk, C. Missing the Obvious: Network Security Monitoring for ICS. Retrieved from <https://www.sans.org/cyber-security-summit/archives/file/summit-archive-1493692381.pdf>

Use of Taps and Span Ports in Cyber Intelligence Applications. Retrieved from <https://support.ixiacom.com/sites/default/files/resources/whitepaper/915-6898-01-taps-span-ports-cyber-intelligence-applications.pdf>

Hoffman, M. (2019, September 19) Gaining Endpoint Log Visibility in ICS Environments. Retrieved from <https://www.csiac.org/journal-article/gaining-endpoint-log-visibility-in-ics-environments/>

Janesko, J.(2018, May 28) Passive Analysis of Process Control Networks. Retrieved from <https://www.sans.org/reading-room/whitepapers/ICS/passive-analysis-process-control-networks-38450>

Hoffman, M. (2019, February 18) Gaining Endpoint Log Visibility in ICS Environments. Retrieved from <https://www.sans.org/reading-room/whitepapers/ICS/gaining-endpoint-log-visibility-ics-environments-38835>

Macaulay, T., & Singer, B. (2016) *Cybersecurity for Industrial Control Systems*. CRC Press

Henning, C. (2009, October 14) Five Criteria for Choosing an Industrial Ethernet. Retrieved from <https://us.profinet.com/five-criteria-for-choosing-an-industrial-ethernet/>