Global Information Assurance Certification Paper

Challenges for IDS/IPS Deployment in Industrial Control Systems

GIAC GCIA Gold Certification

Author: Michael Horkan, mhorkan4223@gmail.com
Advisor: Barbara Filkins
Accepted: July 31, 2015

Abstract

Intrusion Detection and Prevention Systems (IDS/IPS) are a key component of defense-in-depth strategy for information systems. Industrial Control Systems (ICS) and Supervisory Control and Data Acquisition (SCADA) systems need to incorporate this technology in order to properly defend against a growing threat landscape. This paper examines how to deploy this technology in a sample ICS/SCADA setting, identifies hurdles that both industrial control system vendors and asset owners must overcome in order to make IDS/IPS deployment successful, and provides recommendations for both vendors and owners in order to approach the use of these technologies.

This paper is written with two audiences in mind. It is intended for the enterprise IT professional who is familiar with security technologies and best practices, but unfamiliar with ICS/SCADA, as well as ICS/SCADA engineers and managers who lack experience in enterprise security.

1. Introduction

Cyber security in the realms of industrial control systems (ICS) and supervisory control and data acquisition (SCADA) is a developing field. For many years, ICS/SCADA was an area that relied on custom embedded devices and clear-text communications protocols that were not designed with security in mind. Any thought toward their potential vulnerability was answered through a perceived lack of knowledge regarding the technology ("security by obscurity") and a reliance on lack of external connectivity (the "air gap").

The attacks of September 11, 2001 resulted in a new focus on security issues associated with critical infrastructure. The US Department of Homeland Security defines critical infrastructure as "assets, systems, and networks, whether physical or virtual, so vital to the United States that their incapacitation or destruction would have a debilitating effect on security, national economic security, national public health or safety, or any combination thereof" (DHS, October, 2013). The resiliency to penetration and denial-of-service for ICS/SCADA became a new focus for owners, operators, vendors and governments. By 2007, a standard for industrial cyber security, ISA99 (now known as ISA/IEC 62443) emerged (Tofino Security, n.d.).

Key to an effective information security strategy is the concept of defense-in-depth. This strategy demands that several layers of defense be employed in defending a system against malicious attacks. To employ defense-in-depth, an asset owner defines what in their system is most valuable to protect, and employs as many defensive measures around it as possible. This causes an attacker to have to jump through as many "hoops" as possible in order to scan the system, compromise it, steal information and/or damage assets, and avoid detection while doing so.

For example, in the ICS/SCADA realm, a gasoline pipeline asset owner may consider pressure and flow sensors and control valves to be their most valuable control system assets. Defense-in-depth applied to protecting these would include the following:

- Locked physical enclosures
- Physical access portal sensors

Michael Horkan, mhorkan4223@gmail.com

- Security cameras
- Control network segmentation and access control
- Control network monitoring

Historically the perimeter defense for a system is the firewall. Traditional firewalls only guard the gate through filtering IP addresses and ports. What if an attacker finds a way to circumvent the firewall through architecture weaknesses, physical access, or improper employee security habits? What if the attack is carried out through trusted IPs and/or ports, or sensitive data is stolen through these trusted channels?

The means of examining protocols and application-level conversations for signs of both attack and exfiltration – a "watchdog" for data in motion – is found in Intrusion Detection and Prevention Systems.

1.1. The Importance of IDS/IPS in Cyber-Security Solutions

Both IDS and IPS offer a window into events happening in a customer's environment. These systems can alarm in response to and/or block network security or endpoint security events. They are a key component of defense-in-depth for information systems – and yet, adoption of these tools is not pervasive in the industrial control systems market.

According to Doug Wylie of NexDefense, "Anomaly and intrusion detection technologies capable of providing visibility and actionable intelligence about an operating control system continue to be slow to adopt in industry. In comparison to business enterprise IT systems, Industrial Control Systems (ICS) have particular and unique challenges in how they operate and how they have evolved, both of which lend to slow adoption of new network technologies and software solutions that help to detect and protect systems from potential harm" (Wylie, D., personal communication, 6/4/2015).

A primary reason for this is the potential for incorrectly flagging or blocking valid information; these incidents are referred to as "false positives". If the rules that an IDS/IPS uses to define malicious information traffic are improperly defined or simply too broad, packets that are in fact harmless can be categorized by the tool as malicious. In

the case of IDS, this results in a large number of logged records that can make searching for records of true malicious traffic that much more difficult. For IPS, false positives can result in necessary applications having interruptions or full denial-of-service.

It is the latter possibility that would cause the most concern for ICS asset owners and operators. ICS asset owners consider loss of view, the inability to monitor critical control functions, and loss of control, the inability to affect the process, to be the worst-case scenarios. For enterprise IT systems, loss of service can result in business process delays, but rarely impacts revenue directly, assuming proper disaster recovery and business continuance procedures. Such is not the case for industrial control systems; many industries can see literally thousands of dollars in lost revenue per minute of downtime. Further, loss of view or control can be a threat to human and environmental safety; loss of life, damage to equipment and natural resources are real possibilities (Abrams, M., Weiss, J., August 2007)(Dragos Security, December 2014)

Further exacerbating this issue is the possibility of the ICS system being responsible for critical infrastructure services such as electrical generation, distribution, or water supply. A disruption in control to a coolant pump in a nuclear power plant is far more impactful due to its potential to cause harm to life and the environment than a loss of confidential data. Such a concern is what drives government agencies like the US Federal Bureau of Investigations (FBI) and the US Department of Homeland Security (DHS) to take a proactive approach to encouraging ICS vendors and owners toward better security protections.

2. The State of Cyber-Security in Industrial Control Systems

2.1. Off-the-shelf and Connected

Enterprise Information Technology's (IT) primary purpose is the storage, movement, and protection of data. Data can be confidential information, intellectual property, logs and records of personal or commercial nature, and communications. In the industrial sector, similar technology to IT can be purpose-built but for a different focus – manufacturing things, moving people and materials, supervising and controlling utilities.

This is known as Operational Technology (OT). It is the technology that controls, views, and supervises "things in motion."

Compared to enterprise IT, security in Operational Technology is a relatively immature concept. Concerns with OT security have existed for decades, but interest in infrastructure security and resilience peaked in response to two major events – a renewed focus on security following the September 11, 2001 attacks and the analysis of the Stuxnet worm in 2010 (IHS Technology, 2014).

For most of its existence, ICS has utilized proprietary communications hardware and protocols between operator interfaces, engineering workstations, and control processors. These networks and protocols were designed with an implied trust model. This means that components were assumed to only communicate truthful information, properly formatted for the consumer, with no authentication required.

Recent years have seen a shift toward the use of network technologies like Ethernet and the TCP/IP protocol stack. This reduces costs, eases installation and enhances maintenance ability, but also makes the control system architecture more susceptible to vulnerabilities inherent in these popular, well-researched technologies. Legacy systems continue to use inherently insecure protocols such as Telnet, FTP, OPC (OLE for Process Control) and RPC/DCOM.

Many ICS asset owners still trust in the myth of the "air-gap" – the idea that their industrial control network (OT) is physically segregated from the enterprise (IT) network - as their primary strategy of defense.

The air gap myth has been discussed and refuted (Byers, E., August 2013). Suffice to say there are many ways in which air gaps are violated: USB transfers, mobile maintenance or vendor notebook computers that intermittently connect, modems, serial connections, and undocumented network bridges created by employees. Even CDs and DVDs represent a threat vector that breaks the myth.

Still, malicious actors need not get that creative to access a large number of industrial systems. There are a great deal of industrial components connected directly to

the internet. The internet search engine SHODAN (www.shodan.io) reveals several thousand devices that communicate using ICS standard protocols.

Figure 1. Shodan ICS Radar Output of Devices by ICS Protocol (Source: https://ics-radar.shodan.io)

Connections like this may result from a lack of understanding and proper design of system needs. As an example, one ICS asset owner installed both enterprise IT and ICS networks side-by-side yet separate. Maintenance engineers installed an unauthorized bridge between the two networks because of legitimate work needs. A proper network design that had accounted for all valid use cases would have resulted in proper network segmentation and secure controls.

The myth of the air-gap must be abandoned altogether in the interest of advancing the maturity of industrial security. ICS owners must start to embrace the concept of a connected plant floor that supports traditional security safeguards - software updates, security patches, antivirus and IDP/IPS alert rule updates. Secure architectures to support this for ICS/SCADA need to be designed and implemented.

2.2. Examining the Threat

ARC Advisory Group reports that ICS asset owners are roughly equally concerned with vulnerability introduced by connection of OT to enterprise IT as they are to insider

threat (ARC Advisory Group, 2014). This suggests that the ideal placement of IDS/IPS in the control system should be balanced; protection must be supplied at both the boundary (where the control system connects to the enterprise network) and internally to the control network (where a disgruntled engineer or maintenance worker may have physical access). In SCADA systems, where remote transmitter units may be physically separated literally by miles (as in oil, gas, and power distribution networks) physical security is of particular concern. These stations typically operate without the presence of maintenance or operators for extended periods of time, and often lack adequate physical protections and monitoring capability.

Insider threat is not exclusive to deliberately malicious behavior. An engineer who takes his or her laptop home and contracts malware may inadvertently introduce this threat into the work environment. So, too, any vendor, consultant or contractor who introduces unmanaged computers or storage media to the plant environment, either via direct connection, USB, or wireless client. The possibility of intrusion is high, and this further emphasizes that boundary protection alone is inadequate.

Embedded ICS, such as industrial controllers, communication adapters, printers, and barcode scanners, utilize common protocols and expose common services. This makes them just as vulnerable to malicious network traffic as Linux- or Windows-based workstations and servers. However, the direct impact of network-based exploits against such targets is far more likely to be denial-of-service than less-visible and potentially long-lasting exploits like command-and-control (C2) servers (e.g., the receivers for attackers' manual inputs) or pivot points (e.g., network "relays" that attackers use to penetrate further into systems.)

The reasons for this are varied. Lack of common hardware platforms, the limitations of real-time operating systems, minimal available memory and processor bandwidth, and lack of options for persistence (the manner in which exploits survive reboots and power cycles) are some of these. In fact, there is little need for such arbitrary code execution exploits on embedded platforms in the industrial environment, since Windows- and Linux-based platforms are present and much more ideal for this purpose. Stuxnet took

advantage of workstations with Step 7 PLC programming software to edit the controller's user application and inflict damage (Falliere, Murchu, & Chien, 2011).

3. IDS/IPS in Industrial Environments

Intrusion Detection and Prevention Systems come in two major architectural variations – host (or endpoint) based (HIDS/HIPS) and network based (NIDS/NIPS). Each type has its role in secure architecture. For a detailed introduction to these types, see the Appendix.

3.1. A Sample Architecture for IDS/IPS Deployment

Figure 2 presents a simplified, version of Cisco and Rockwell Automation's Converged Plantwide Ethernet architecture (CPwE). This architecture is not intended to be representative; it is presented for illustration of concepts alone.

The Enterprise IT network is separated from a Demilitarized Zone (DMZ) by a firewall. The DMZ contains a Web Proxy Server, a Remote Desktop Gateway, and an FTPS File Server. Each of these acts as a go-between for the Enterprise and Manufacturing Zones. The Manufacturing Zone is separated from the DMZ by the firewall, and contains all assets necessary for the control system's operation – Programmable Logic Controllers (PLCs), Human-Machine Interface (HMI) workstations, maintenance and engineering computers, sensors, input and output (I/O) adapters, drives, and supporting application servers. Routers and switches make up the network infrastructure.

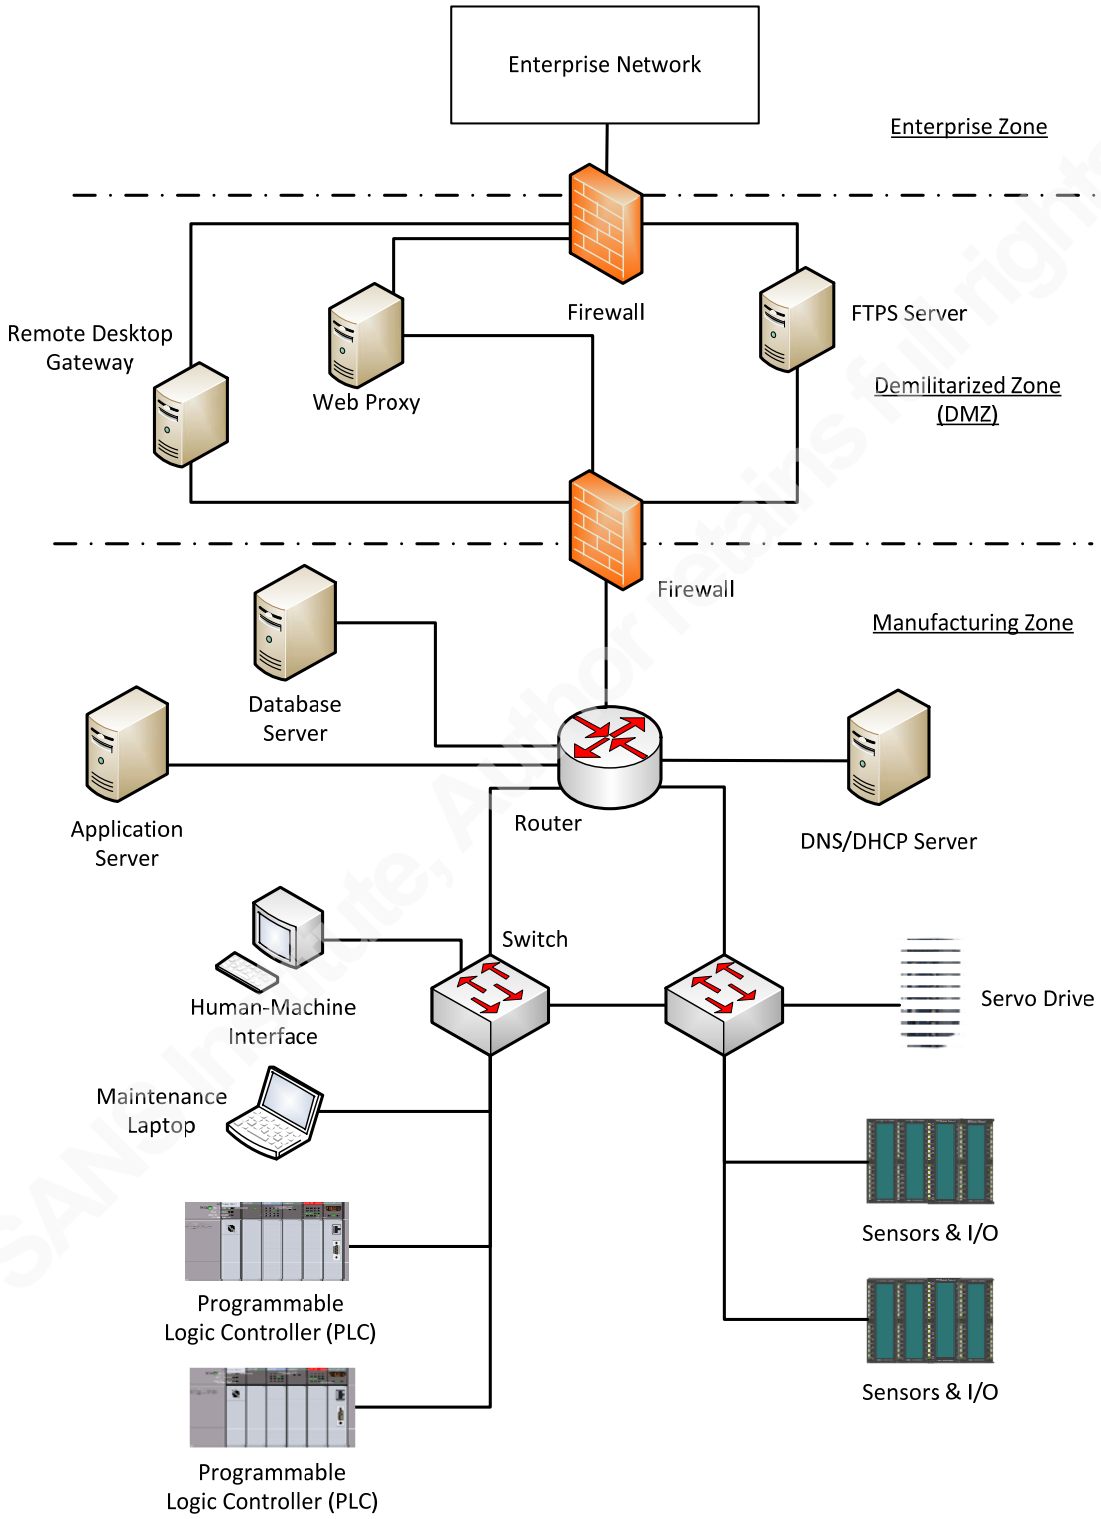

Figure 2. Example Simplified CPwE Architecture

The elements in the above architecture may utilize many application protocols that are common to enterprise IT, as well as exclusive OT protocols such as Ethernet/IP, Common Industrial Protocol (CIP), Modbus, DNP3 and many others that are capable of being transported over TCP/IP. An industrial IDS/IPS must be capable of parsing and examining all of these for threats.

3.2. ICS Threat Scenarios

Scenario 1 – Enterprise network-sourced threat

If a compromise occurs at the enterprise IT level, the means for an attack to reach the ICS is through the DMZ. The various proxy servers located in the DMZ are typically hardened and specialized for a particular purpose – these are known as "Bastion Hosts" (Dillard, n.d.).

Proper configuration of firewalls and use of a DMZ means that all connections terminate in the DMZ. There are no direct connections from the Enterprise to the Manufacturing Zone and vice-versa. A proxy server is a dual-homed host that accepts connections from both the Enterprise and Manufacturing zones, and passes application protocols between them.

Figure 3 illustrates HTTPS traffic passing between the Enterprise and Manufacturing Zones using a web proxy. The web proxy requires its own signed TLS certificate to establish two HTTPS sessions.

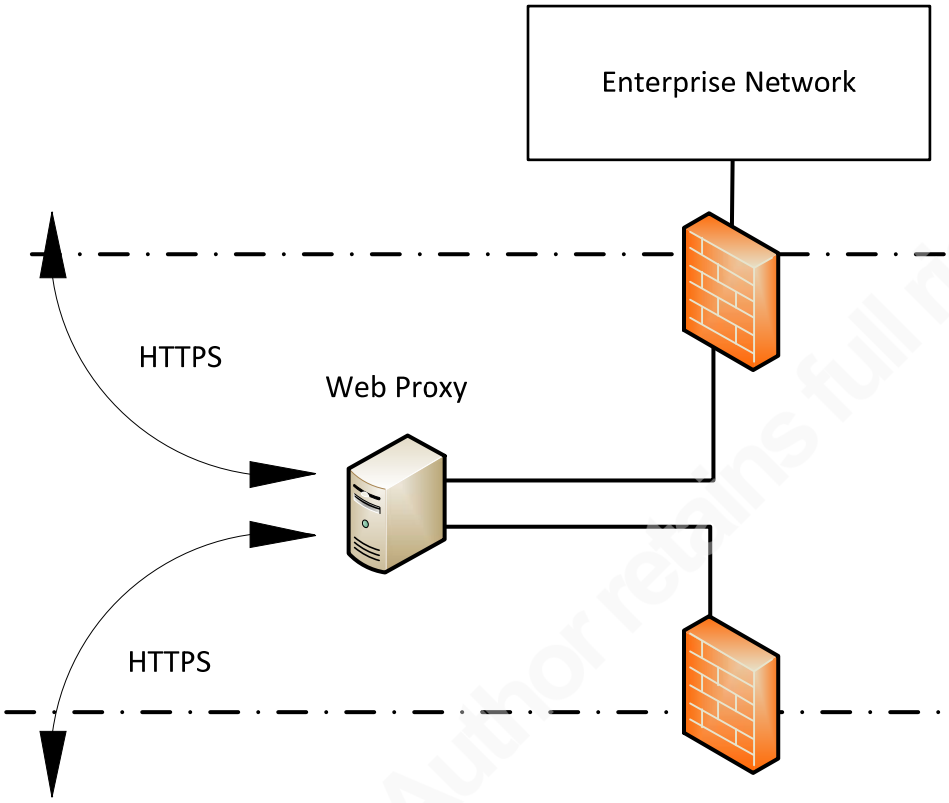

Figure 3. Web Proxy with HTTPS traffic in the DMZ

Here the attacker may attempt to compromise the server itself, but not necessarily as a more "elegant" method may be to simply encapsulate malware through protocols that the server is dutifully handling. IDS/IPS should be employed to cover all possibilities – HIDS/HIPS to detect on-host compromise of the proxy server itself, NIDS/NIPS to detect compromise attempts directed at the proxy server at the network connection, and a host-internal NIDS/NIPS to inspect the application traffic that the proxy is passing between the zones. Figure 4 illustrates this deployment.

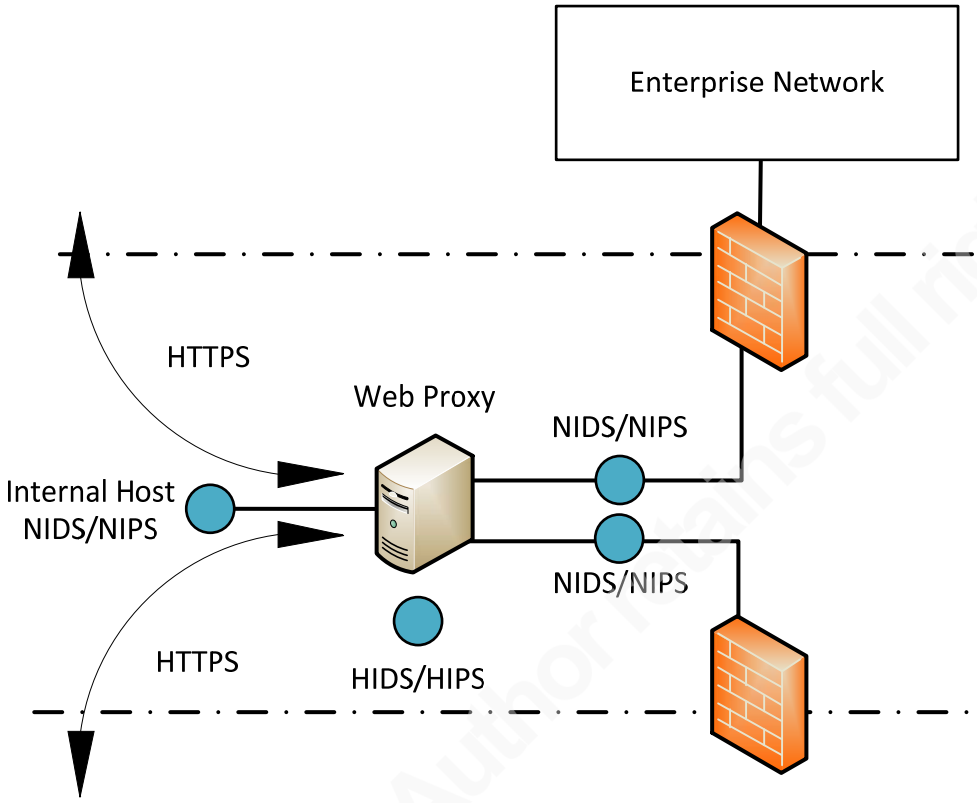

Figure 4. Proxy server sensor deployment

In the above figure, the right-most NIDS/NIPS sensors may be contained in the firewalls. Firewalls that encapsulate traditional IP and port filtering capabilities with NIDS/NIPS and application control capabilities are referred to as Next-Generation Firewalls.

In the DMZ, sensors would be tuned to detect exploits only in the protocols the proxies are configured to pass. In our example, these would be HTTPS, Remote Desktop (RDP), and secure file transfer (FTPS). All other protocols would be blocked, and detected attempts to send them to the DMZ would be flagged as alerts.

Scenario 2 – Insider Threat

In this case, we concern ourselves with malware introduced directly at the plant floor (in the Manufacturing Zone). Suppose an unscrupulous maintenance technician uses his or her work laptop to surf the Internet at home and the laptop becomes infected with malware. Subsequently, the technician connects the laptop to the Manufacturing Zone and introduces the malware to the ICS. Another possible source of malware at the

Manufacturing Zone is contractors, who often require and supply their own connections to the internet. A scenario like this occurred at the Ohio Davis-Besse nuclear power plant, when a contractor connected a T1 line to the plant's corporate network, bypassing security controls that isolated it from the Internet, and introduced the Slammer SQL worm to the plant (Kesler, 2011). In this particular case, the corporate and control networks were not isolated from each other, providing an example of the potential threat vector that contractors can pose to a protected network.

Host-based protection on the laptop, Manufacturing Zone workstations, and application servers is certainly a good idea here and a necessary first line-of-defense. In the event that the HIDS/HIPS fails, a network-based sensor could detect IP/port scans, Nmap-style OS and service fingerprinting, as well as industrial protocol-based malware traffic. Stuxnet was known to have established a peer-based update mechanism on a LAN (Falliere, Murchu, & Chien, 2011). In that case, network monitoring would have been useful to detect the worm. Monitoring such as this would serve to protect any attached hosts that are susceptible to integrity compromise (such as other workstations or human-machine interfaces) as well as embedded devices that are more vulnerable to denial-of-service attacks (such as programmable logic controllers).

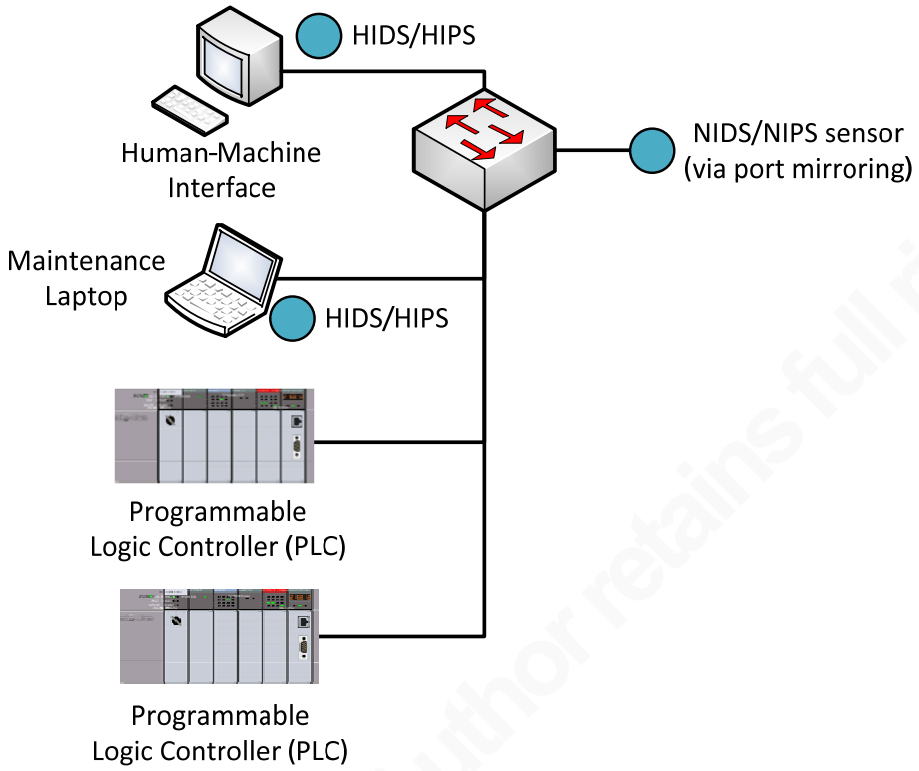

Figure 5. Sensor placement in the Manufacturing Zone of the Sample Architecture

Threats on the inside of the control system network are particularly concerning in SCADA systems, where remote terminal units (RTUs) are often installed at distances measured in miles from a central control room. Consider Figure 6, a representation of a SCADA architecture. A central control room is used to supervise multiple RTUs.

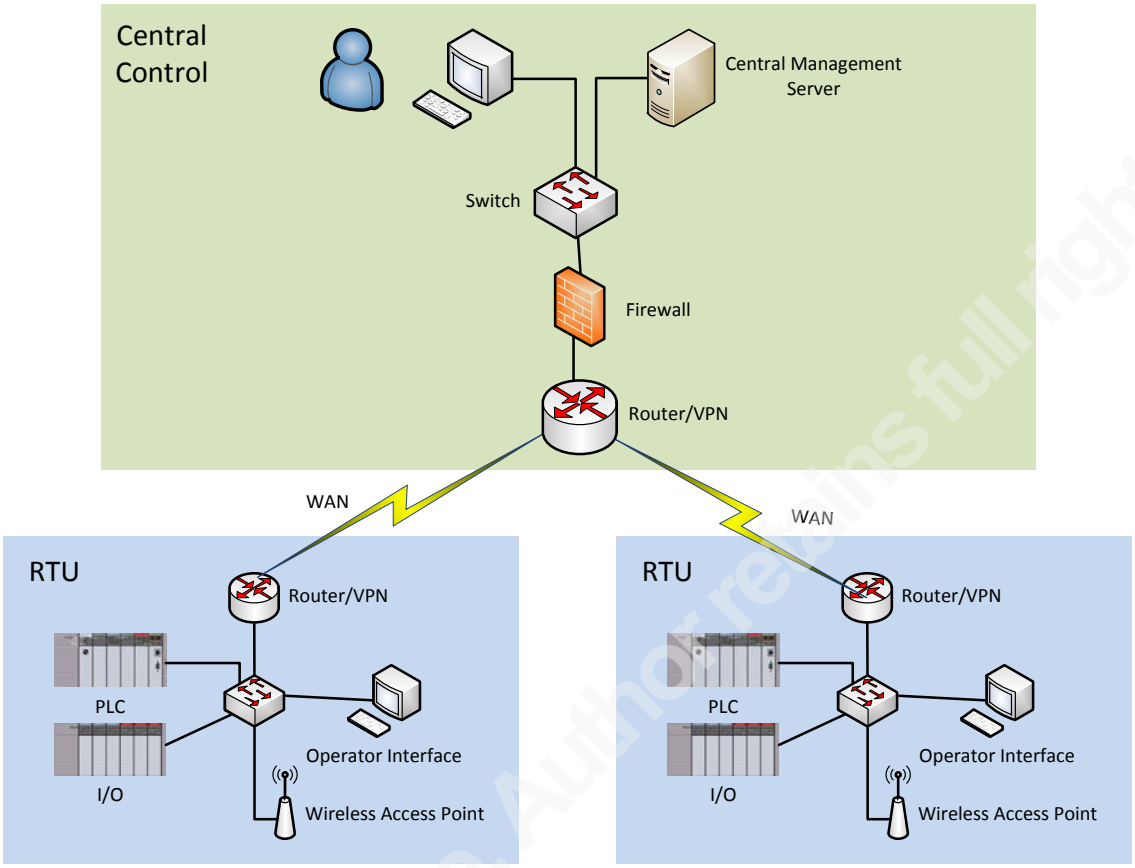

Figure 6. Sample SCADA Architecture

If an attacker defeats the physical security of an RTU, or alternatively, compromises a wireless access point internal to the RTU, he or she could pivot to other RTUs across the system. The attacker could establish a command-and-control network that encompasses all RTUs, invisible to the central control center.

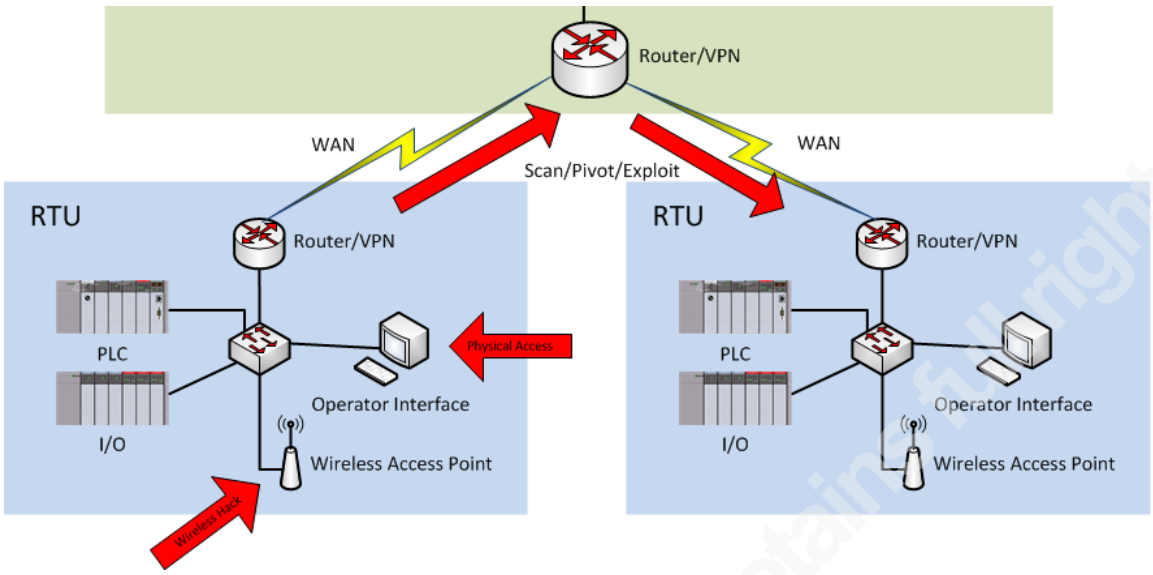

Figure 7. Peer RTU Compromise in Sample SCADA

Network-based IDS/IPS, deployed to monitor the central control network RTU interface at a minimum, would combat this threat.

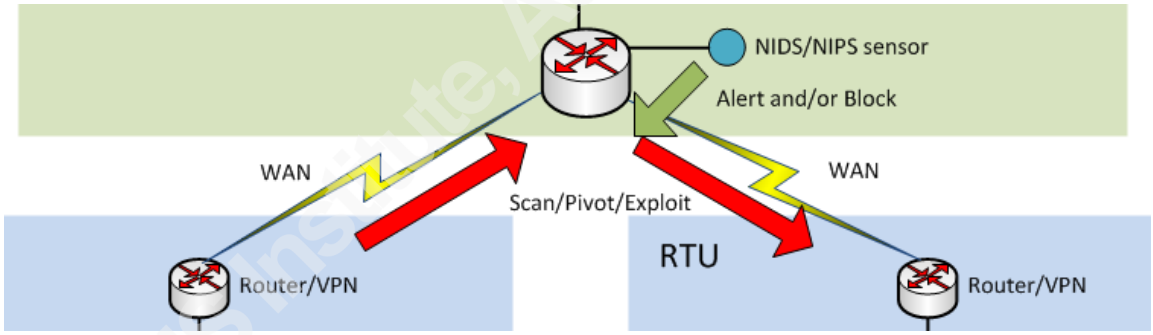

Figure 8. NIDS/NIPS Cross-RTU Protection

3.3. The Challenges for ICS Vendors

ICS asset owners often view security as a component to purchase for their control systems; they want to buy components that make their control system "secure", are not obtrusive to runtime production, can operate in an automated fashion and do not require a large investment in training and maintenance. This is counterintuitive to the trained security professional, who understands that security is not something to "bolt on", but must be "baked in" as part of the secure development lifecycle of components,

architecture reviews and adjustments, human training and behavioral adjustments, and continuous monitoring with frequent oversight.

So the ICS vendor is faced with managing customer expectations and training customers to understand that making ICS secure cannot simply be bought. There is no "magic bullet" they can purchase that will make an ICS without proper zone segmentation, without an effective update plan, without proper access controls, without concept of least-privilege, and without best practices for operators, engineers and maintenance personnel have a better security profile.

What follows is a set of recommendations for both vendors and asset owners to make ICS security better with the incorporation of IDS and IPS technologies and processes.

3.4. Recommendations for ICS Vendors

1. Provide a simplified functionality

Cyber security knowledge and staffing are limited in ICS environments, and many asset owners find it difficult to build and maintain this competency (ARC Advisory Group, 2013). Given the relatively low instance of IT technology experience on the plant-floor, ICS vendors will have to focus on functionality that is easier to understand and integrate.

The complexities of an IDS, including SNORT rules, regular expressions, and command-line interfaces, may be par for the course in security and network technologies (indeed, their interface impenetrability may be considered even desirable to security professionals), but owners of operational technology in the ICS domain will not necessarily see the value of this complexity.

2. Provide pre-processors for industrial protocols

IDS and IPS pre-processors must be designed that understand ICS/SCADA protocols. These preprocessors will make rule creation and coverage both simpler and more robust.

Consider the CIP Stop command packet generated by Metasploit's "auxiliary/admin/scada/multi_cip_command" module. Its purpose is to cause a denial-of-

service by telling an industrial controller to shut down. Figure 9 shows a breakdown of the EthernetIP/CIP payload.

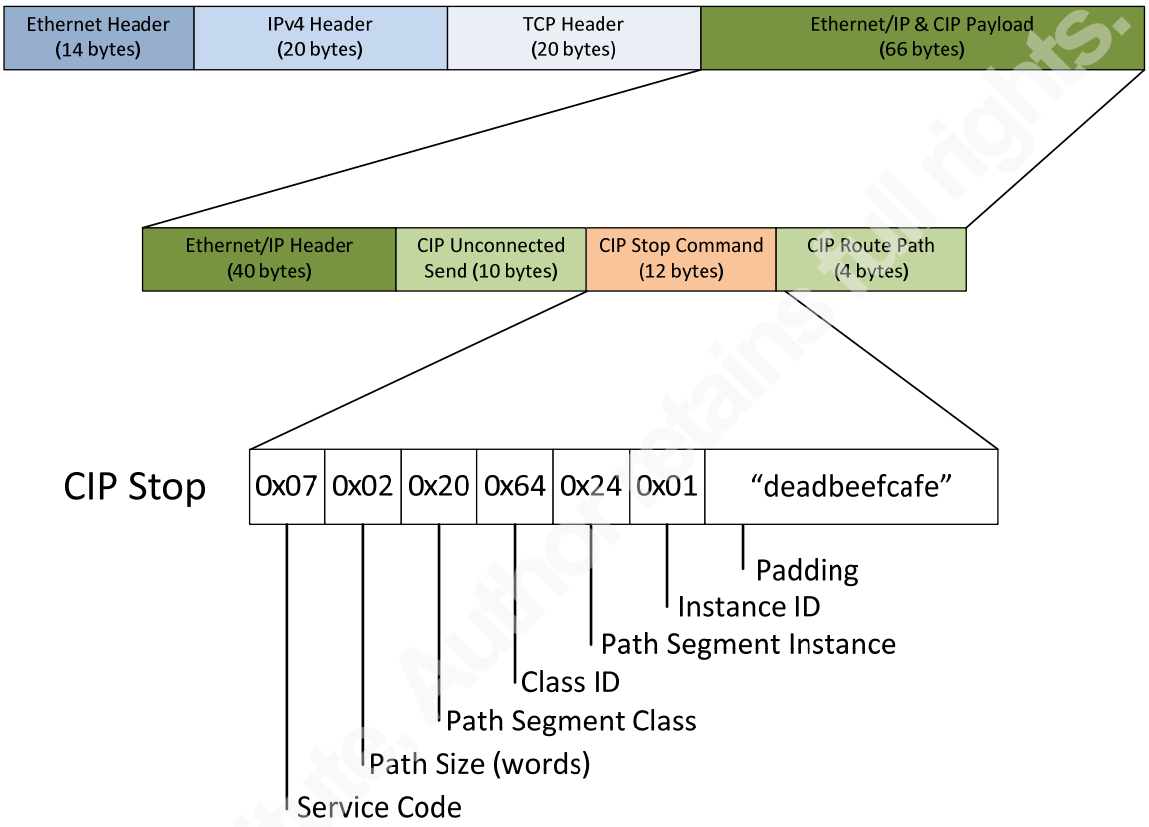

Figure 9. Packet structure of Metasploit CIP Stop Command exploit

This particular packet could be detected and stopped by a simple Snort rule such as this:

```
alert TCP $EXTERNAL_NET any -> $HOME_NET 44818 (msg:"caught CIP STOP exploit!"; flow: established, to_server; content:"|070220642401|"; offset:50; depth:6; sid:70000000;)
```

The rule triggers off of the CIP Service Code through Instance ID hex string, at a TCP payload offset of 50 bytes. We specify these rule attributes to be as specific as possible, in order to avoid false positives. This rule is not adequately robust, however, since the CIP protocol allows for multiple CIP service commands to be encapsulated into a single packet:

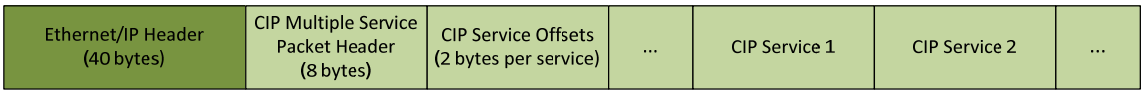

Figure 10. Multiple CIP service encapsulation

Using this valid encapsulation, the CIP Stop command could be placed anywhere within a given packet that contains multiple services, rendering the Snort rule above useless. The rule would be evaded due to the variable offset within the TCP payload.

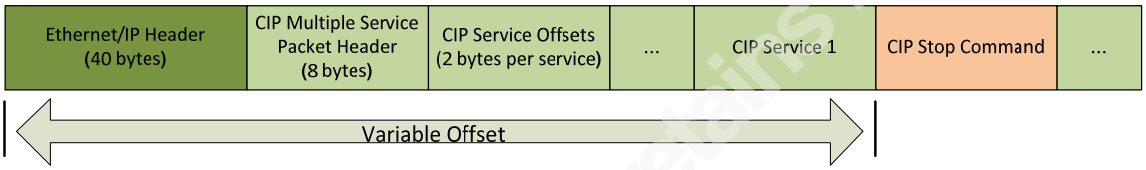

Figure 11. CIP Stop Command Snort rule evasion

This is where a pre-processor, programmed with the protocol intelligence to decode CIP, would allow a more robust rule to be created that could catch an IDS/IPS evasion attempt like this.

SNORT currently supports only two industrial protocol pre-processors: DNP3 and Modbus. These are common protocols for the power distribution industry. Further pre-processors will be necessary for other IP-based industrial protocols like Common Industrial Protocol (CIP), Profinet, EtherCAT, SERCOS III, Powerlink and others.

3. Supply an Industrial SIEM

A Security Information and Event Management system, or SIEM, is an important piece of what makes IDS and IPS workable. This is a supervisory system that receives and aggregates signals from many IDS/IPS sensors and allows asset owners to view and filter results in report form.

A SIEM also aggregates system events related to security: logins/logoffs, privilege escalation, and application alerts are some examples. For industrial systems, events like scheduled downtime, material shortage, shift change and others could be correlated with security events for security incident detection.

Suppose, for example, that program changes to an industrial controller are never legitimately performed outside of scheduled weekend downtime. A SIEM with awareness of scheduled downtime could flag any program configuration traffic detected outside of this boundary as a potentially malicious event, and then annunciated that event with priority.

3.5. Recommendations for ICS Asset Owners

1. Employ Robust Network Design, Tailored to the Process

Without proper network design, monitoring with IDS/IPS is of limited usefulness. If a network is under-designed, the possibility of rogue connections that bypass critical security controls goes up.

Networks in industrial environments are often under-designed for what is required of them. Ensuring that components can communicate is the only consideration. Here is where the lack of convergence between IT and OT is directly and demonstrably impactful. Engineers whose primary area of expertise is process control are expected to "wear the hat" of a network architect, and only occasionally. This fails to give proper credit to an important area of expertise.

Networks must be designed to consider all of the following:

- Resilience: will physical operation or safety be affected by the loss of a switch, router, cable or communication adapter?
- Authorized use-cases: how will operators, engineers, managers and contractors use the network? Are all valid use cases considered and mitigated?
- Operational segmentation: separation of operational units (work cells), input/output traffic, supervisory traffic, synchronous versus asynchronous traffic, controller-to-controller peer messaging, secure tunneling and others.
- Load Balancing: analysis of necessary throughput (packets-per-second), latency, overhead, datagram size effects.

- Physical security: how physically accessible are network infrastructure components, controllers, I/O devices (discrete and analog I/O, drives, sensors)

Design considerations like these affect an industrial system's resistance to attack or accident, prevent information leakage, and facilitate separation of network protocols. All of these benefits are good prerequisites to the next recommendation.

2. Use the "P" only for Non-Critical Traffic

The possibility of false positives is a real concern as long as this technology is in its infancy in the ICS domain. Organizations that currently utilize network intrusion detection technology are generally averse to using it as a prevention method (in other words, as a NIPS) for fear of blocking critical traffic (ARC Advisory Group, 2014). Until industrial protocol parsing becomes ubiquitous in IDS/IPS, and as long as a better definition of what is "normal" traffic in this domain remains elusive, use of this technology in IPS form presents a risk of operational downtime. A deserved, or perceived, reputation as the cause of operational downtime will kill the adoption of this technology - quicker than anything else.

For this reason, ICS owners should apply this technology in Protection Mode (blocking malicious traffic) only where traffic is not directly critical to control. Figure 12 is an illustration of critical versus non-critical traffic in the manufacturing zone of the representative architecture.

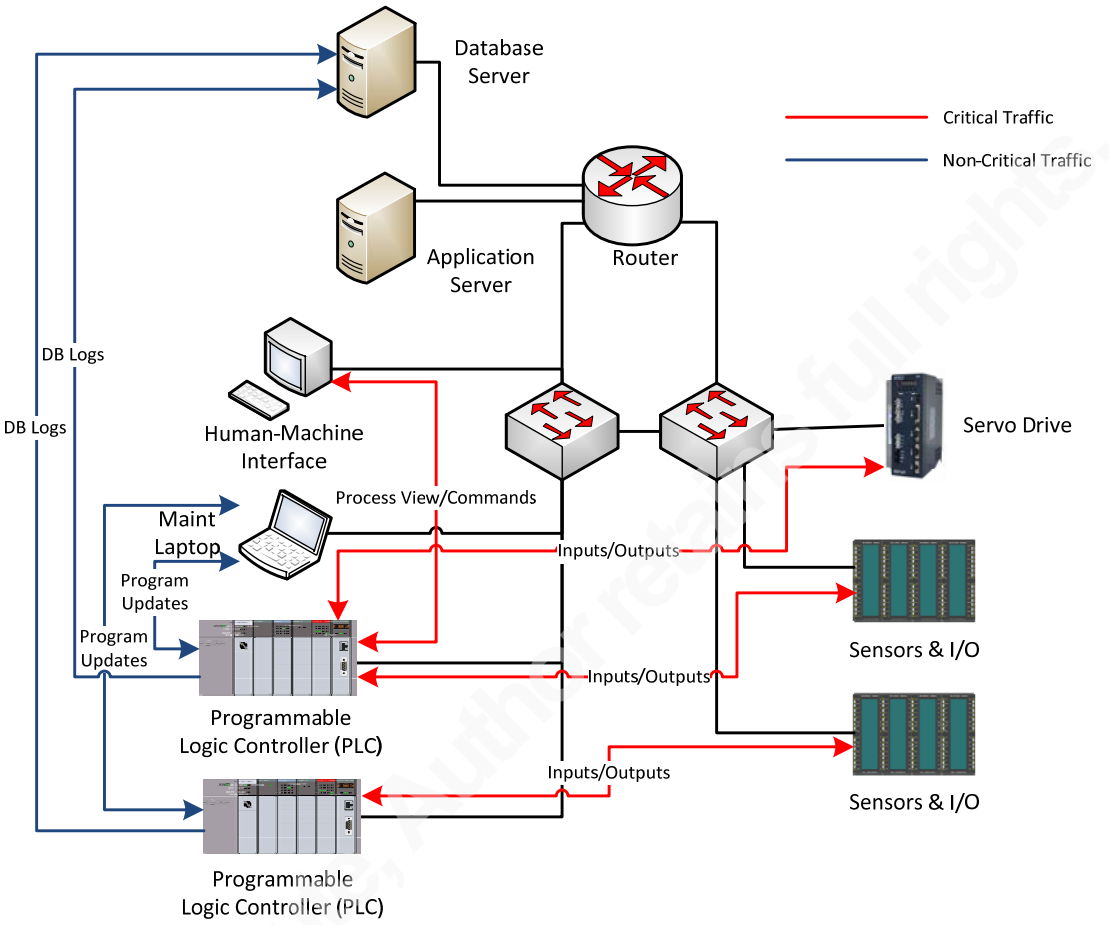

Figure 12 Critical versus Non-Critical Traffic

Understand that this breakdown is extremely dependent on the process and the owner's needs, and must be evaluated on a case-by-case basis. For example, database logs may be as critical to an asset owner as the physical control of a process itself; an owner may determine that they are willing to shut a process down rather than let it continue without this information.

A side benefit of taking a detection-only strategy is the ability to learn more about an owner's ICS network traffic – in terms of both what is normal and in possibly discovering unexpected connections. For example, an owner who holds both an expectation of air gapping as well as proper network usage would be well-informed if an IDS uncovered Youtube traffic destined for an ICS-internal host.

3. Take advantage of the static nature of ICS

One aspect of ICS/SCADA that could make it ideal for IDS/IPS deployment is its fundamentally static, deterministic nature. Once IP addresses and protocols are established within a system at commissioning, they rarely change (upgrades excepted). According to the Department of Homeland Security:

Network traffic between ICS components remains mostly invariable for extended periods of time. An HMI polls specific data points from a PLC at regular intervals. Well-tuned alarms remain relatively "quiet" under normal operating conditions. Sensors and I/O serve up data chunks of the same size at configured polling rates. It is this determinism that makes a network anomaly such as in intrusion stand out, to be caught and annunciated by an intrusion detection system.

This is a network anomaly approach to defense that can be very successful in the ICS/SCADA realm – as opposed to enterprise IT where the host and protocol pool, IP addressing, and network traffic profile are all far more dynamic.

4. Conclusions and Further Considerations

ICS vendors should come out of the gate with products that are competitive in performance and feature set with their enterprise IT counterparts. Industrial Automation product vendors should take advantage of the experience of companies like Cisco, Tipping Point, Juniper and IBM and utilize partnerships to more quickly take advantage of a rising opportunity. A survey of ICS/SCADA owners, vendors and consultants showed that the percentage of industrial companies using HIDS/HIPS will grow by 20% over the next three years, and the use of Security Information and Event Management (SIEM) by nearly 30% (ARC Advisory Group, 2014)

For an enterprise IT vendor, the benefit of this partnership would be in co-branding with a name that ICS owners and operators trust. Currently, Rockwell Automation, Siemens, ABB, and Honeywell are among several well-established and trusted names on the plant floor bill-of-material.

ICS/SCADA system vendors and asset owners have significant work to do in order to improve industrial security. IDS and IPS can and should play an important part in this work, once a good baseline security profile is established. Vendors can take the

lead in this effort, not only in the products they sell but in the knowledge they impart to their customers. ICS asset owners have displayed both a desire and a willingness to let their control system vendors be their guides on technology, best practices, and regulatory compliance. Vendors embrace this role as the "control experts" and must leverage that reliance to raise the bar on industrial security. Together, both groups must actively engage governing agencies to ensure effective security practices are represented in current and future regulations.

Government authorities need to invest in updating their recommendations as well as regulations in light of newer threats and newer strategies for defense. The Department of Homeland Security's recommendations for better securing industrial control systems, referenced in this document, are six years old as of this writing (DHS, October, 2009).

An effective IDS/IPS deployment will require the ability to update alert signatures regularly. Inability to update will make this technology ineffective. This is what makes industrial systems' update management, or more accurately lack thereof, a problem for continuous monitoring. System updates are often actively discouraged in ICS/SCADA. This can be due to a concern for downtime, or due to strict regulations like those of the FDA that require extensive re-testing when system updates are made.

There are technological, organizational, and personnel skillset drivers that keep IT and OT apart (Harp, D., & Gregory-Brown, B., April 2015). For IDS/IPS to succeed in industrial environments, these areas will need to converge. Effective deployment if IDS/IPS in this market segment will require plant floor capability to install, maintain, and update sensors and central aggregation servers. To meet this need, owners will be required to engage human resources with the skills necessary to understand and apply this technology.

To conclude, we quote the ARC Advisory Group: "Given the complexity of today's ICS systems, the growing cyber threat landscape, and the multitude of products available, developing an effective ICS cyber security strategy requires considerable expertise. This and the need for constant system maintenance have spawned a large market for cyber security services." (ARC Advisory Group, 2013)

5. Appendix – Types of IDS/IPS

Host-Based Intrusion Detection/Prevention (HIDS/HIPS)

Host-Based Intrusion Detection/Prevention (HIDS/HIPS) is endpoint software that monitors for malicious activity, and can alarm (HIDS) and possibly block (HIPS) operations on the endpoint client that are perceived as malicious. Often, this functionality is included endpoint protection suites that bundle it with AV (antivirus), firewall, and application whitelisting. HIDS/HIPS can monitor system memory, the file system, and removable media.

The host-based solution allows for a more detailed view of malicious software's effect on an endpoint; however, it should not be relied on exclusively since any intrusion that goes undetected has the potential to compromise the detection software itself. Architecturally, the host-based solution allows examination of network packets that the host network adapter can see, but ONLY those that the host can see.

In industrial systems, host-based IDS/IPS can be deployed on human-machine interfaces, workstations, and servers. Embedded devices like programmable logic controllers and communication adapters may eventually incorporate this technology once processor speeds and memory are available to support it.

Figure A-1 is a simplified architectural representation of Host-Based IDS/IPS. Malicious traffic, both inbound (ingress) and outbound (egress) is represented by the red arrows, detected and/or prevented at the host.

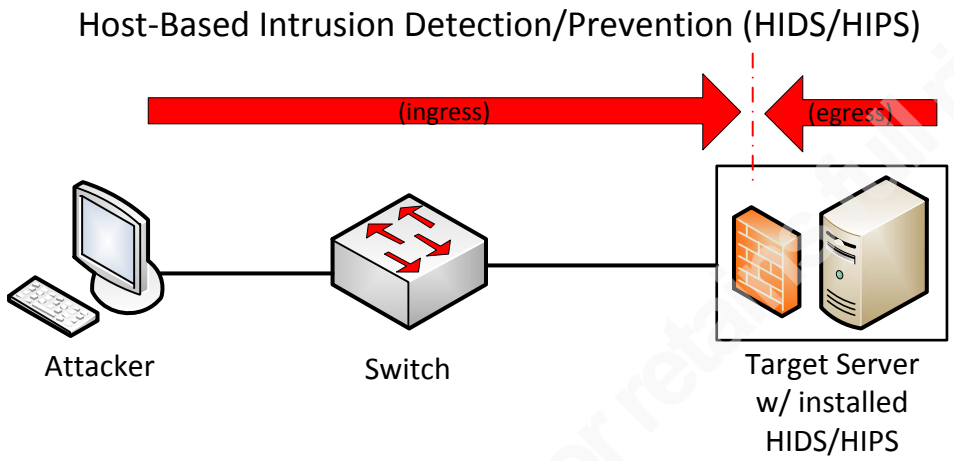

Figure A-1. Host-Based IDS/IPS. Red arrows indicate malicious traffic

Network-Based Intrusion Detection/Prevention (NIDS/NIPS)

Network-based intrusion detection and prevention is a technology that examines traffic between network components and either alarms (NIDS) or blocks (NIPS) anomalous or malicious network packets on the wire.

This solution requires an appliance – in the form of a dedicated endpoint – positioned in the architecture to read network packets at one or more locations. This appliance can connect with a network tap or a switch port that mirrors traffic.

This form of intrusion detection lacks the ability to see signs of intrusion within endpoints, but casts the wider net of examining packets from several endpoints near-simultaneously, given proper sensor placement. Nearly every intrusion must entail generating network traffic in order to be useful to attackers, either through transmitted exploits or command-and-control (C2) packets.

Industrial systems could implement network-based IDS/IPS similarly to how it is done in enterprise IT – through span ports on network switches or through network taps.

Figure A-2 is a simple representation of network-based IDS/IPS, implemented with a network tap, protecting a target server.

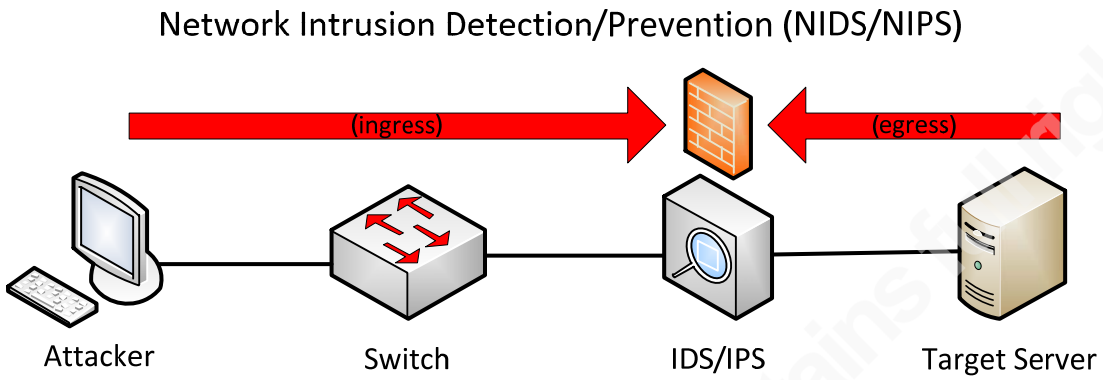

Figure A-2. Network-Based IDS/IPS. Red arrows indicate malicious traffic

References

Abrams, M., & Weiss, J. (August 2007). *Bellingham, Washington Control System Cyber Security Case Study*. http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Bellingham_Case_Study_report%2020Sep071.pdf

ARC Advisory Group (2013). *ICS Cyber Security Global Market Research Study, Market Analysis and Forecast Through 2017.* Retrieved through subscription

ARC Advisory Group (May 2014). *The Real State of Industrial Cyber Security.* Retrieved through subscription

Byers, E. (August, 2013). *The Air Gap: SCADA's Enduring Security Myth*. Communications of the ACM, Volume 56, No. 8

Dillard, K. (n.d.). *Intrusion Detection FAQ: What is a Bastion Host?*. http://www.sans.org/security-resources/idfaq/bastion.php

Dragos Security (December 2014). *ICS Cyber Attack on German Steelworks Facility and Lessons Learned*. https://dragossecurity.com/blog/9-ics-cyber-attack-on-german-steelworks-facility-and-lessons-learned

Falliere, N., Murchu, L., & Chien, E. (February, 2011). *W32.Stuxnet Dossier, Version 1.4.* http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf

Harp, D., & Gregory-Brown, B. (April 2015). *IT/OT Convergence: Bridging the Divide*. http://ics.sans.org/media/IT-OT-Convergence-NexDefense-Whitepaper.pdf

IHS Technology (June 2014). *Cybersecurity in process and discrete automation.* Retrieved through subscription

Kesler, B., (2011). *The Vulnerability of Nuclear Facilities to Cyber Attack.* Strategic Insights, Spring 2011, http://edocs.nps.edu/npspubs/institutional/newsletters/strategic%20insight/2011/SI-v10-I1_Kesler.pdf

Michael Horkan, mhorkan4223@gmail.com

Tofino Security (n.d.). *ISA/IEC 62443 Standards*. https://www.tofinosecurity.com/why/isa-iec-62443

US Department of Homeland Security (October, 2009). *Recommended Practice: Improving Industrial Control Systems Cybersecurity with Defense-In-Depth Strategies.* https://ics-cert.us-cert.gov/sites/default/files/recommended_practices/Defense_in_Depth_Oct09.pdf

US Department of Homeland Security (October, 2013). *What is Critical Infrastructure*. http://www.dhs.gov/what-critical-infrastructure