DRAGOS

Key Considerations for Selecting an Industrial Cybersecurity Solution for Asset Identification, Threat Detection, and Response

Introduction

In the industrial cybersecurity industry today, there are two main challenges facing those responsible for protecting their organizations' networks: a largely unknown threat landscape and a shortage of experienced personnel with effective threat detection and response capabilities.

With the number of threats to the ICS industry increasing every year, organizations are assessing their need for industrial cybersecurity in response to these challenges, and many are looking for support from ICS cybersecurity platforms that can make the most of their budgets, supplement their defenses, and give them more resilient security postures.

With numerous ICS solutions and approaches available, it can be difficult to determine which one is right for your organization. This paper outlines considerations organizations should take into account when researching ICS cybersecurity platforms, based on the Dragos team's first-hand experience of what's needed for effective industrial cybersecurity.

Understanding how different capabilities drive different outcomes, this paper will help you consider your organization's specific requirements and assess the best solution for your industrial operations.

Asset Identification

You can't protect assets you don't know you have. The ability to continually understand and have visibility into your environment should be considered a prerequisite step to prevention, detection, and response. Asset identification provides the ability to know and visualize what assets your industrial environment is comprised of. Asset identification should not be confused with asset management though. Asset management could include internal inspection of control systems as well as patch management of those assets.

When assessing asset identification capabilities of an industrial cybersecurity platform, it is important to consider:

Visual Representation

A visual representation of the environment in an asset map extends asset identification from a static spreadsheet to a dynamic understanding of not only the assets in the environment, but how those assets communicate, their relationships with other assets, and their involvement in the industrial processes in which they function.

Visualizing this asset data in a map intuitively gives insights into how the environment is *designed* and how it *operates*. The ability to view and navigate industrial assets in a visual representation significantly lowers the barrier to entry and reduces the resources necessary for proactive and defensive cybersecurity--with an objective to minimize any kind of downtime associated with cybersecurity.

Sorting and understanding this asset map is important but should not be difficult. Utilizing auto-zoning capabilities reduces the time spent classifying and grouping assets and makes asset data readily available for analysis immediately after deployment, avoiding “bake-in” time.

Passive vs. Active Scanning

In industrial environments, safety of operations is always mission one. Passive analysis is not intrusive, does not involve the logistical challenge of deploying and managing agents, and it prevents any potential impact to operations. More importantly, passive analysis allows defenders to see network traffic at near real time without the fear of disruption or harm to the environment. Additionally, passive analysis supports threat detection by providing the content to do deep packet analysis and traffic analysis; this is otherwise impossible with active scanning.

Active scanning involves interrogation and probing attempts which can cause inadvertent denial of service and downtime to devices or industrial processes. Significant effort is being made in this space to provide more reliable and robust ways to perform active asset identification in safer approaches; however, many of the methods still require knowledge of what your assets are ahead of time. Additionally, the information gathered is not as useful for other use-cases such as threat detection.[1] However, this active scanning can be useful for deeper vulnerability identification against known vulnerabilities. Given the current state of industrial vulnerabilities and the potential risk to operations active scanning is often best reserved for lab and test environments.

Changes Over Time

Asset identification in industrial cybersecurity must be able to track changes over time and enable historical timelines of each assets' status to be able to assess if a change is intentional (non-malicious), or unintentional (accidental or malicious). This tracking can be done in multiple ways. One common method is to create snapshots, or baselines, of the environment to compare to later. This is useful to identify changes but requires accurate and consistent snapshots. Another method is to have a complete historical timeline that is dynamic to compare at different points of time. This method is more useful for use-cases that support investigations and incident response.

Comprehensive, Contextual Environmental Perspective

An industrial cybersecurity platform's asset identification and threat detection capabilities can be amplified by leveraging multiple data sources to give a more comprehensive view of the environment and the incidents that have occurred within it. Inputs should not rely solely on endpoint agents and associated data, but instead, offer flexibility and an array of approaches and inputs. The ability to ingest multiple data sources creates a more comprehensive and contextual representation of risk.

Network traffic is a significant and important data source that has a large return on investment, as one sensor can gain visibility across an entire network--thereby, reducing deployment costs while generating a comprehensive view. A defining feature for this category of offerings is the ability to monitor network traffic and perform deep packet inspection of industrial protocols to extract out an understanding of the systems and their communications.

Though network traffic analysis will help with detection and help generate a scope of an incident, it alone does not have the depth of data to provide high-fidelity correlations of data or explain what

1 Reference: Dragos 2017 YIR Vulnerability Report
https://www.dragos.com/media/2017-Review-Industrial-Control-Vulnerabilities.pdf

happened to a host or device. Achieving root cause analysis of incidents requires more than just network data. Other good data sources to be able to collect and correlate include firewall logs, host-based logs from operations technology (OT) systems, syslog from industrial components and controllers, and information and alarms from data historians.

Expertise in Industrial Environments, Devices, and Protocols

Industrial Environment

In-depth understanding and knowledge of industrial environments, collectively, is the first step to better threat detection and response; however, an in-depth understanding of individual industrial environments is absolutely critical to effective defense (e.g., understanding the environment of and differences between a power plant owned by Company A; oil refineries 1, 2, and 3 owned by Company B; or a manufacturing plant owned by Company C).

The highly heterogeneous nature of industrial environments means there are no cookie cutter approaches to cybersecurity, as there are at least slight differences between every company (and thus every adversary targeting that company).

Opportunistic Attacks in Industrial Environments

There are increasing examples of automated propagation of worms, viruses, and ransomware used in industrial environments that lead to shut down, denial of visibility, and long recovery times and costs. The best engineered industrial processes are increasingly reliant on strong cybersecurity to prevent impact from opportunistic attacks ranging from worms and data destroying ransomware such as NotPetya.

Targeted Attacks in Industrial Environments

For adversaries to create a specific impact to the industrial process, they must be well-versed in the specifics of the target's environment. For example, in an oil refinery, an adversary may need to know a range of details depending on their objectives:

- The temperature at which they could trigger a chemical reaction (or lack thereof)
- What safety and protection safeguards are deployed or engineered into the systems
- What systems monitor normal operations
- How the equipment is designed and what opportunities the design reveals for malicious intentions
- What could cause a system to stop functioning properly

Cybersecurity approaches that negatively impact the industrial environment create a business continuity risk--thus, an industrial cybersecurity solution should understand the organization's cybersecurity architecture and product functionality, identify malicious attacks, and analyze how adversaries progress through an attack. This understanding of threats, ranging from opportunistic ransomware to sophisticated attackers, should be continually updated with the latest intelligence to identify threat behaviors and learn from their patterns.

Industrial Devices

Industrial devices are designed to operate for long periods of time without intervention. They are often designed to be minimal in terms of needs for physical power, compute power (processing), and complexity to maintain 24x7 uptime in often harsh environments or remote locations. This inherent "simplicity" means industrial devices are not built to be particularly resilient to a broad spectrum of system inputs.

Industrial devices are also often deployed for decades at a time, far longer than modern-day electronic devices. This means that these "simple" industrial devices will remain in place and unchanged, even as modern technology progresses--especially in situations where updating or upgrading creates more risk and downtime compared to the benefits of the update or upgrade.

An industrial cybersecurity solution must be able to operate within the bounds of the capabilities of the industrial devices it protects. For example, a PLC may "freeze" if it receives an unknown signal. In such a case, a cybersecurity solution utilizing active scanning would have detrimental effects. Instead of freezing, it may be designed to restart automatically. Or perhaps it's part of a network of PLCs that is programmed for a daisy-chain set of reactions. Being able to identify the type of device by their specific purpose and design is a critical requirement for an effective industrial cybersecurity solution.

Context and understanding of these industrial devices is also important. An industrial cybersecurity solution should have an understanding of the different types, roles, and relationships that exist in the environment. This means not only should the solution understand what role a human machine interface plays, but also how it's different from an engineering workstation, and how those different roles can influence the potential consequence of an attack. Finally, different types of devices have a different exposure and, therefore, have different attacks associated with them. An industrial cybersecurity solution should understand these differences and focus on attacks that are most relevant to the various types of devices.

Industrial Protocols and Deep Packet Inspection

Communication within an industrial network is what brings the entire system to life. An industrial power plant, manufacturing facility, or processing plant contains a vast number of interconnected devices working together; but the diversity of vendors and lack of standardization results in a myriad of ICS protocols being transmitted and received between different vendor equipment, as well as between the same vendor's equipment. Since control signals to devices are sent via these protocols, an industrial cybersecurity solution must be able to understand and interpret the meaning and impact of all the different protocols to prevent blind spots in your security approach.

To understand protocols and the inherent valuable data within them, deep packet inspection (DPI) is necessary for comprehensive insight into the device communications occurring on your networks. DPI allows you to dig through the layers of data in a specific packet to get to the data that actually matters, such as the application information and protocol headers, which would allow you to see if a device is communicating in a way it shouldn't. If a solution doesn't offer DPI, you don't get complete visibility of device communication, which means suspicious--or malicious activity--that could be a potential vulnerability or threat gets overlooked. DPI provides a deeper layer of device communications, so you can identify threats quicker and more comprehensively.

Contact Dragos at info@dragos.com for the current list of supported protocols.

Beyond ICS protocols

It's an oversimplification to only rely on an understanding of ICS-centric protocols. Many attack vectors observed recently surround IT-centric protocols, such as Microsoft's Server Message Block (SMB) used for windows-to-windows file sharing, HTTP protocol used for web pages, secure shell (SSH), and others. Being focused only on the ICS protocol "layer" could create visibility gaps in your security posture.

In industrial environments, there is more to consider besides ICS protocols; understanding function codes and identification of their abuse is important in identifying and preventing threats. The ability of DPI for passive identification of systems and software such as engineering workstations, HMIs, historians, and the suite of applications in DCS or SCADA environments provides valuable contextual information to assess appropriate responses to threats.

Efficacy and Efficiency in Detecting Threats

Reduce Dwell Time and Minimize Downtime

The ability to reduce dwell time is a key evaluation criterion for selecting an industrial cybersecurity solution. One deterministic way to reduce dwell time is to evaluate if the solution is able to identify and detect threats in an environment on day 1, as soon as it has been deployed (i.e., no "bake-in" period). This enables immediate notifications and alerts to gain critical insight into the activity on your networks and to detect any threats that exist.

If the solution requires a bake-in period, you could potentially miss something malicious already occurring on your network--meaning, slower detection and response time, potential risks to operations in general, and impacts to your networks. Additionally, those solutions that require effort to "bake-in" have to be constantly tuned and updated. This activity often encourages analysts to "whitelist" communications and commands that are legitimate. However, adversaries commonly use legitimate protocols, commands, and services to conduct accounts. Poor tuning of solutions requiring a "bake-in" can quickly result in missing attacks and incidents.

Alleviate Alert Fatigue with Efficient Threat Detection

A common problem facing ICS analysts today is alert fatigue; receiving hundreds to tens of thousands of notifications a day about alerts that offer no context as to which are malicious causes, ineffective threat detection, and inefficient response. Because security analysts are a resource pool in high-demand and short supply, maximizing the efficiency of these highly talented professionals is a key evaluation criterion for an industrial cybersecurity platform.

Security analysts should be able to access context-rich notifications, because they offer visibility into the alert itself, the severity of threat level, and provide context of how analysts should respond.

There are several approaches to maximizing the productivity of security analysts, such as machine learning, anomaly detection, and threat behavioral analytics. Each have value in different environments with different requirements.

Machine Learning

Machine learning is a type of modeling which can be a useful and effective option for highly homogeneous environments or inputs. The challenge is that machine learning requires a large dataset to create the appropriate models and must be trained. Training should not just include environmental characteristics but also malicious actions. Industrial environments, while more complicated than commonly thought, are still wildly too small and/or consist of heterogeneous datasets to create highly trained and useful models. Machine learning can be valuable, but it is overall minimal when compared to its applicability in IT counterpart environments. Additionally, because the models are mostly trained on the environment, and not threats, the alerts they generate are numerous and have no context about threats. Common types of machine learning application include protocol behavior analytics and device behavior analytics. Here, machine learning is used to build a profile of normal behavior for specific types of protocols or devices. A more common type on the market is anomaly detection.

Anomaly Detection

Anomaly detection is another type of modeling which results in the identification of rare or abnormal events. If something is misconfigured, or changed when it should not be, anomaly detection is useful. Anomaly detection typically depends on machine learning to sift through large amounts of data that is not possible for human analysts to process.

Anomaly detection can cause some undesirable side-effects, such as a higher number of notifications and false positives and can lead to alert fatigue if utilized as a primary source of detection. Additionally, anomaly detection requires a baseline to compare the new activity against. While creating this baseline, a solution can inadvertently create blind spots by incorporating suspicious or malicious activity into that “normal” baseline. This can quickly lead to increased risks, as many industrial threat actors utilize ‘living off the land’ methodologies and techniques.

The efficiency of threat detection through anomaly detection should also be considered. The process of threat detection based on anomaly detection means something is identified as being “abnormal” without a clear assessment of criticality, impact, or maliciousness. Once detected, an anomaly needs to be investigated to assess whether it is a real threat or not. This requires industrial system expertise to be able to evaluate and can be a fairly long process to gather all the relevant data to make a determination. Unless an industrial security team is readily accessible, the mean time to respond (MTTR) for the organization will increase--which, from a risk measurement perspective is a negative result. If an enterprise is motivating security teams to reduce risk by measuring the number of tickets closed, this further fuels the problem and results in a poor security outcome.

Threat Behaviors

While detecting on anomalies or signature matching can be a useful tool, primary detection should be based on the tradecraft used by known threats for more efficient threat detection. To pinpoint malicious behavior and relieve alert fatigue, threat behavior-based approaches are more effective and efficient.

Threat behavior analytics identify cybersecurity threats through complex patterns of adversary techniques, tactics, and procedures (i.e., they are developed by practitioners who have already hunted and characterized them). The ability to chain together a series of events into a custom

analytic allows for customization of behaviors that have a stronger relevance and fidelity of information, rather than relying on one static, atomic value. In essence, analytics are questions that are asked of your data; because the analyst knows the question and why it's important then the answer back is more quickly utilized as opposed to anomalies or machine learning based approaches which generate numerous answers without knowing the question. As an example, an analytic may be search data for the pattern of a previously observed tradecraft of an adversary such as the downloading of a file to an HMI, followed by reaching out to the internet, followed by interaction of the adversary to operate the HMI.

Threat behavior analytics provide context to the alert when the analytic alerts (e.g., what's happening, why you should care, and what you should do about it) - reducing the number of hours the analyst would spend trying to figure out what's going on. Processing all available data and providing context to alerts prevents analyst fatigue and allows resources to be directed to activity of concern, given the specific environment.

Threat behavior analytics also enable an industrial cybersecurity solution to leverage a "know what you're looking for" approach versus a "notice something odd, and then take the time to figure out if it's something you care about" approach of anomaly detection. An industrial cybersecurity solution that utilizes threat behavior analytics results in a lower false positive rate and increased efficiency.

Indicators

Static indicators are not well-suited as a detection mechanism but are relied upon in information sharing channels to quickly scope for near identical incidents. The ability to ingest and sweep for these indicators across historical data and alert in the future should be a routine and easy-to-use option in an industrial cybersecurity solution. This ability becomes even more valuable when these sweeps are done across a wide range of data sources to include network traffic, authentication logs, application logs, host operating system logs, firewall logs, etc. This allows for more confidence and coverage, but also enables sweeping for different types of indicators, such as file hashes, usernames, and other artifacts beyond simple IP address or attacker infrastructure.

Indicators are not a durable form of threat detection, but they can be useful for quick scoping and forensics post detection.

More information on the Four Types of Threat Detection is available here:
https://dragos.com/media/The_Four_Types%20of_Threat_Detection.pdf

Reduce the Mean Time to Respond

Threat Hunting

It is important to recognize that attackers are often ahead of the defenses in an organization. Threat hunting is necessary to identify the gaps in defenses and to focuses the humans in these areas to find adversary activity. Threat hunting is always a human-assisted endeavor; however, enabling large-scale collection of data from the environment that can be used by a threat hunting analyst for correlation and identification of threat behaviors that aren't yet codified into an automatic detection schema is extremely valuable when identifying an attack in early stages.

Threat hunting refers to applying the human analyst until the technology can catch up. This means easily providing the human, not just with data, but with data that can be easily processed, indexed, correlated, filtered, and searched in novel ways.

Read more about Collection Management Frameworks – Looking Beyond Asset Inventories in Preparation for and Response to Cyber Threats here:
https://dragos.com/media/CMF_For_ICS.pdf

Comprehensive Data Collection and Depth of Visibility

Just as the ability to ingest multiple data sources can amplify asset identification, it also maximizes threat visibility. An industrial cybersecurity platform should be able to ingest a wide variety of data sources--data from asset identification information, pcaps, logs (System, Event, PLC, RTU), historians, and more. Correlating and compiling disparate sources of data in one place for an analyst to search, along with access to raw and processed data, significantly accelerates the threat hunting process.

Ingesting logs that are available in the environment can offer the multifaceted view that's required for analysts to both *detect* and *understand* what has occurred during incidents, along with the ramifications of the activity. The ability to passively ingest, process, and correlate logs in the environment allows for a higher fidelity behavior detection mechanism, identifies blind spots in the environment, and provides the analysts a more comprehensive view to identify context and accuracy of events.

Empower Analysts with Accelerated Ability to Respond

Constant and passive monitoring of the operational environment brings visibility to assets and network communications into a single area for analysis, but what should be expected of organizations for investigation and response?

Without knowledge and experience, response to security threats can be difficult and overwhelming. Checklists, predefined search queries, and case management are often effective tools in assisting analysts in answering the question of "How do I respond?" Otherwise, it is extremely ineffective to investigate an alert without context or to understand the context after impact. Without experience, how to collect, interpret, and respond to data is unknown.

Technologies that provide context around the event or alert, leading the analyst on the correct path, greatly reduce the time to act and increases effectiveness of the response. Solutions that only provide event information (e.g., a new system communication is detected) risk causing the analyst to generate more questions than streamlining the discovery of pertinent answers.

Context amongst the chaos of an event or incident is key. Knowledge transfer embedded within solutions is preferred, reducing the need for the analyst to access multiple disparate systems for information. The source of the context or intelligence is also something to consider, given federal, state, or other regulatory requirements. **Intelligence-driven context matched against system events and notifications can quickly empower a junior analyst to fulfill the requirements of a much senior response role**.

Once an analyst knows what actually transpired to trigger the alert or notification, that data can be used to prevent the event from occurring again, whether that be an adjustment in process, systems,

policy, etc., or a cull on response to a suspected breach. Documenting this into action plans (playbooks) makes this a scalable, repeatable process to enable defenders of all levels to confidently navigate through the investigative and/or threat hunting process by independently following a check list of guidelines authored by senior practitioners.

When evaluating an industrial cybersecurity offering, choose evaluation criteria that match your requirements. If the objective is to reduce the mean time to respond, consciously consider what impacts that; are people able to be more productive, can processes be accelerated or alerts be reduced?

Flexibility, Scalability and Operational Relevance

Flexible Deployment Options

Industrial environments have wide-ranging characteristics; some are controlled environments, while others are highly-variable, remote locations with extreme environmental conditions. At the same time, there are numerous options for industrial cybersecurity connectivity internal to the OT network and externally to the IT network. Deployment criteria of an industrial cybersecurity solution to consider include:

- What kind of systems (historians, devices) can be monitored?
- What types of data can be collected and analyzed?
- Will they be connected to or via the cloud?
- What are the environmental requirements?
- Are there existing investments or toolsets which need to be integrated with? (SIEMs, management infrastructure, ticketing infrastructure)
- Are there any limitations by vendor? (e.g., control system vendor specifications or warranty requirements)

Choosing an industrial cybersecurity platform that can meet your different environments and architectures enables more comprehensive threat visibility than a patchwork of different solutions.

Scalability

Industrial environments have a large number of devices and systems producing a vast volume of data. When selecting an industrial cybersecurity platform, it is important to evaluate whether it can scale to support asset identification of the hundreds or thousands of devices on your network, as well as the gigabytes and, quite possibly, terabytes of data being generated at any given time.

In order to minimize risk visibility and downtime, it is recommended that you are able to capture, process, analyze, and store terabytes of data. This is because the typical ICS adversary takes a year or more to achieve their goal, and historical data is needed to enable accurate analysis.

The rate of data processing is also of importance to reduce dwell time. If it takes days or months to process and analyze industrial network data in order to investigate a threat, there will be a increased dwell time for the adversary to initiate and execute their intentions.

Operational Relevancy

Environments and adversaries are constantly changing. To maintain the value of an industrial cybersecurity solution, it must be continually updated with the latest capabilities and threat detection capabilities. The asset identification map needs to be current and up-to-date to be useful. Behavior analytics need to be updated to detect new adversaries and changes in tradecraft of existing adversaries. If using anomaly detection, the system will need to constantly be tuned. Security teams need to continuously expand their knowledge and experience through updated and expanded investigation playbooks. Keeping up with the threats via regular updates, rather than waiting on product level updates or upgrades, ensure continuity of defensive capability and minimize security risk.

Conclusion

When determining the appropriate industrial cybersecurity solution for your organization, there are many capabilities to weigh and measure value and effectiveness. We recommend creating a checklist based off your specific organization's cybersecurity needs, so you can comprehensively examine your environment, determine budget limitations, and identify areas of weaknesses. Our mission at Dragos is, and will continue to be, to empower the ICS community with in-depth knowledge and independence to determine the most effective ICS defenses to reinforce your organization's security posture, effectively pinpoint threats on your networks, and help you protect and respond to attacks effectively and efficiently.

Please feel free to reach out to us at info@dragos.com for any questions you may have.

Useful References

A Dragos Industrial Control System Security Reading List:
https://www.dragos.com/blog/20181213-TOC-Reading-List.html

Industrial Control System Threats:
https://www.dragos.com/media/2017-Review-Industrial-Control-System-Threats.pdf

Collection Management Frameworks – Beyond Asset Inventories for Preparing for and Responding to Cyber Threats:
https://www.dragos.com/whitepapers/CMF_For_ICS.html

How to Respond to Industrial Intrusions:
https://www.slideshare.net/DragosInc/how-to-respond-to-industrial-intrusions

The Four Types of Threat Detection with Case-Studies in Industrial Control Systems (ICS):
https://www.dragos.com/blog/FourTypesOfThreatDetection.html

Dragos 2017 YIR Vulnerability Report:
https://www.dragos.com/media/2017-Review-Industrial-Control-Vulnerabilities.pdf