AUTOMATION 2020

Volume 2 | April

Industrial Cybersecurity

- An expanding attack surface puts plants at risk
- The next big step for EtherNet/IP
- Bridging the IT/OT cybersecurity divide
- Understanding USB interface threats
- Microsegmentation and edge security

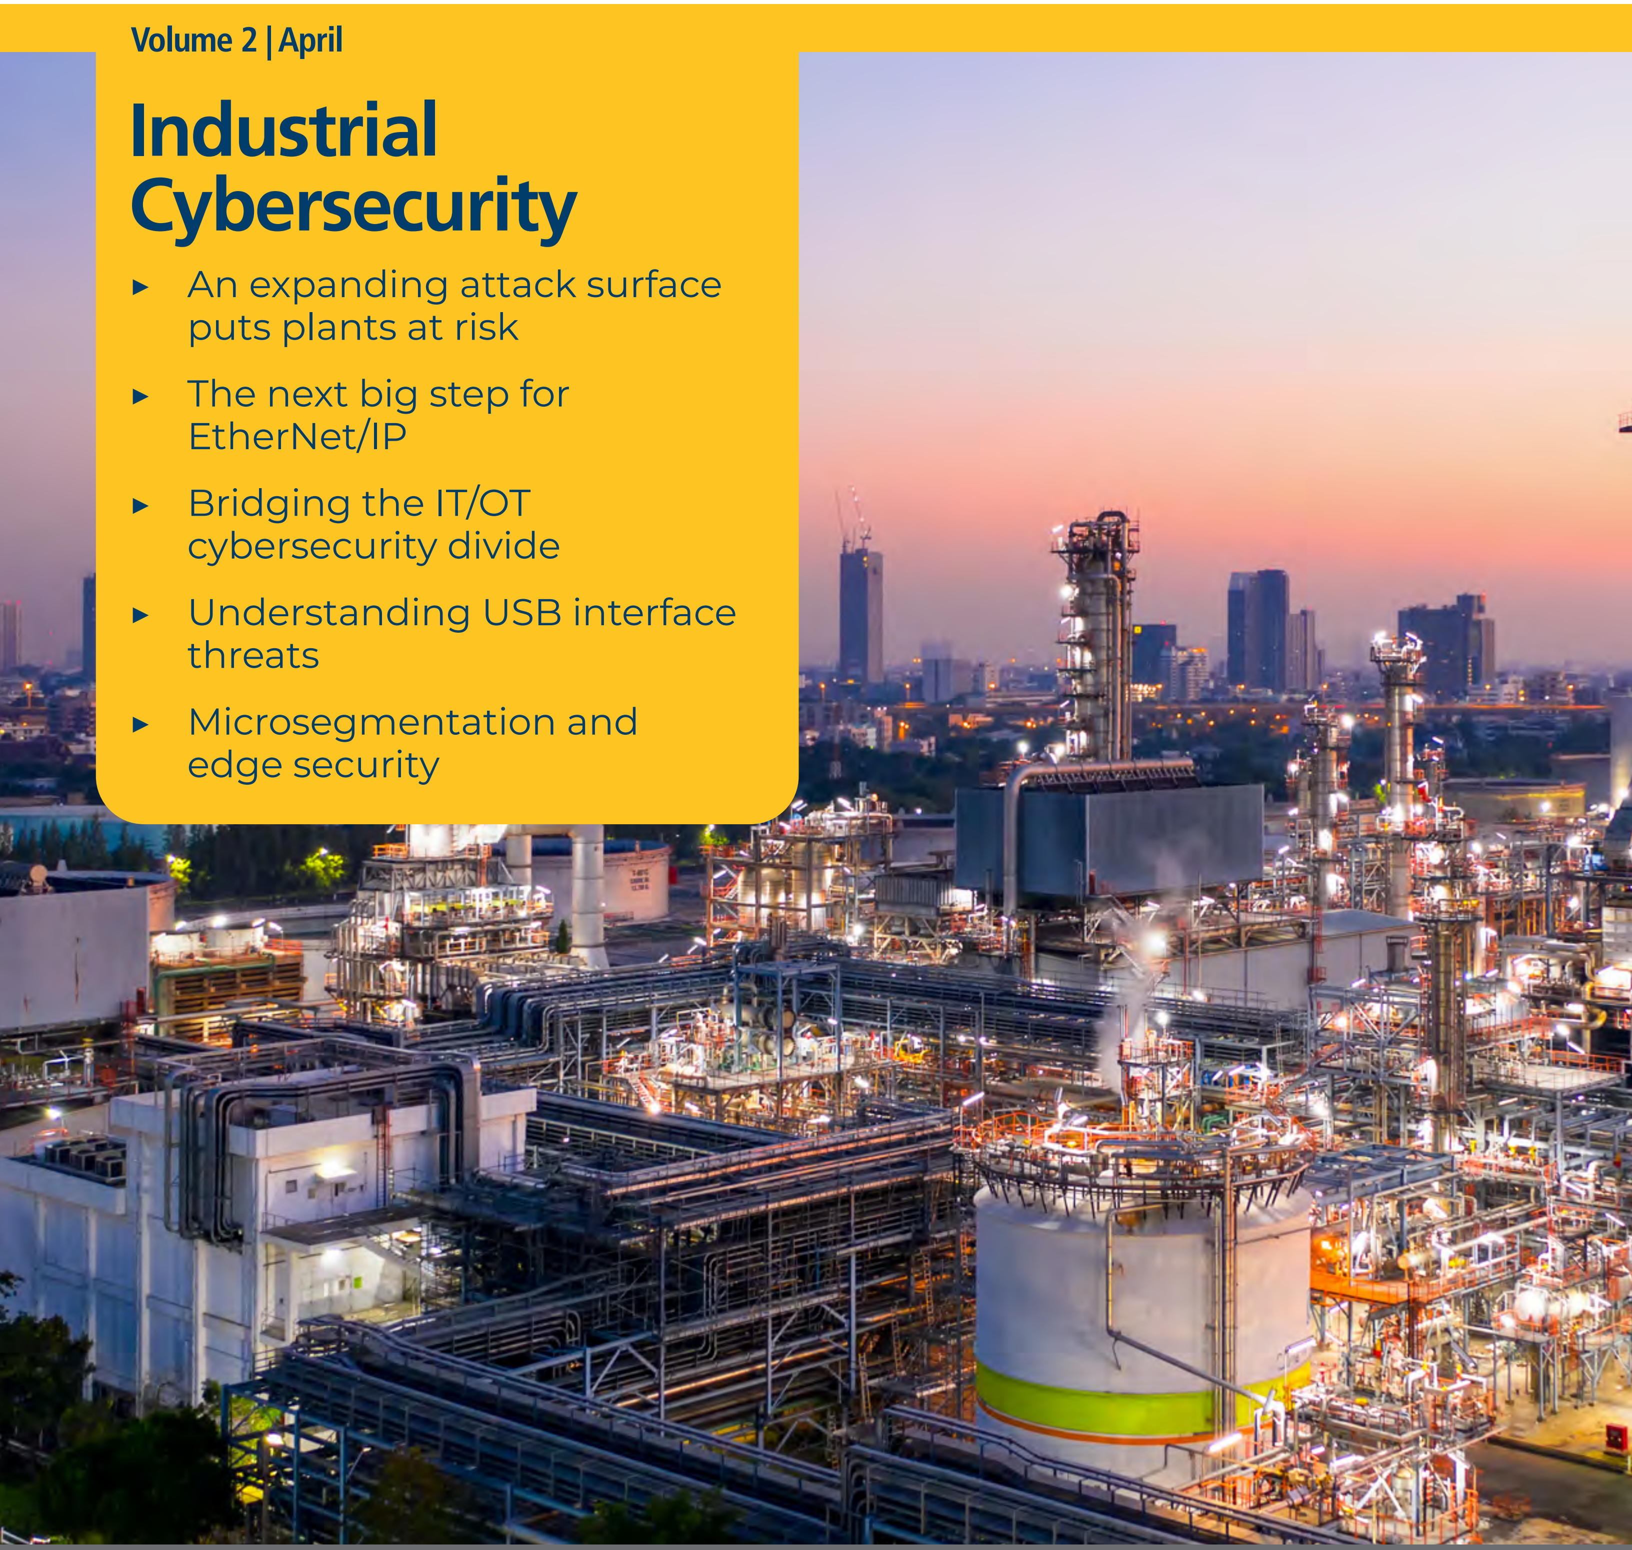

AUTOMATION.COM

A subsidiary of the International Society of Automation

Introduction

Industrial Cybersecurity

Information technology (IT) and operational technology (OT) departments approach operations in very different ways. While innovation, upgrades, and a global view are common in IT, industrial OT organizations are mainly focused on getting a system to run reliably with as few changes as possible. So, it makes sense that IT and OT approaches to cybersecurity are different as well. IT environments were always connected to a network and the Internet, so security risks have been top of mind for decades. OT departments, on the other hand, have relied on “air gaps” to separate industrial networks from the rest of the world. With digitalization and ubiquitous connectivity, IT and OT have converged. And there is a lot of cybersecurity learning that needs to be shared.

In this edition of **AUTOMATION 2020** from Automation.com, find out how to bolster business resiliency in the face of cyberthreats, how to protect against threats to industrial control systems, the latest on USB port intrusion threats, the next big step in EtherNet/IP security, and more. Formerly called *Advancing Automation*, this ebook series includes sponsored and curated articles from a variety of experts on industrial cybersecurity. Subscribe online to not miss an issue of these guides to best practices and cutting-edge insight for automation professionals.

Table of Contents

AUTOMATION 2020 VOL 2 | APRIL
CYBERSECURITY

HONEYWELL FORGE | Cybersecurity
EMEA
Sites
ASSET DISCOVERY
3,065
ASSET ALARMS
SOFTWARE PATCH & AV
SECURE
25
3,653
1,033
721
493
470
348

Back to USB School

USB attacks blur the line between a network threat and a local physical threat

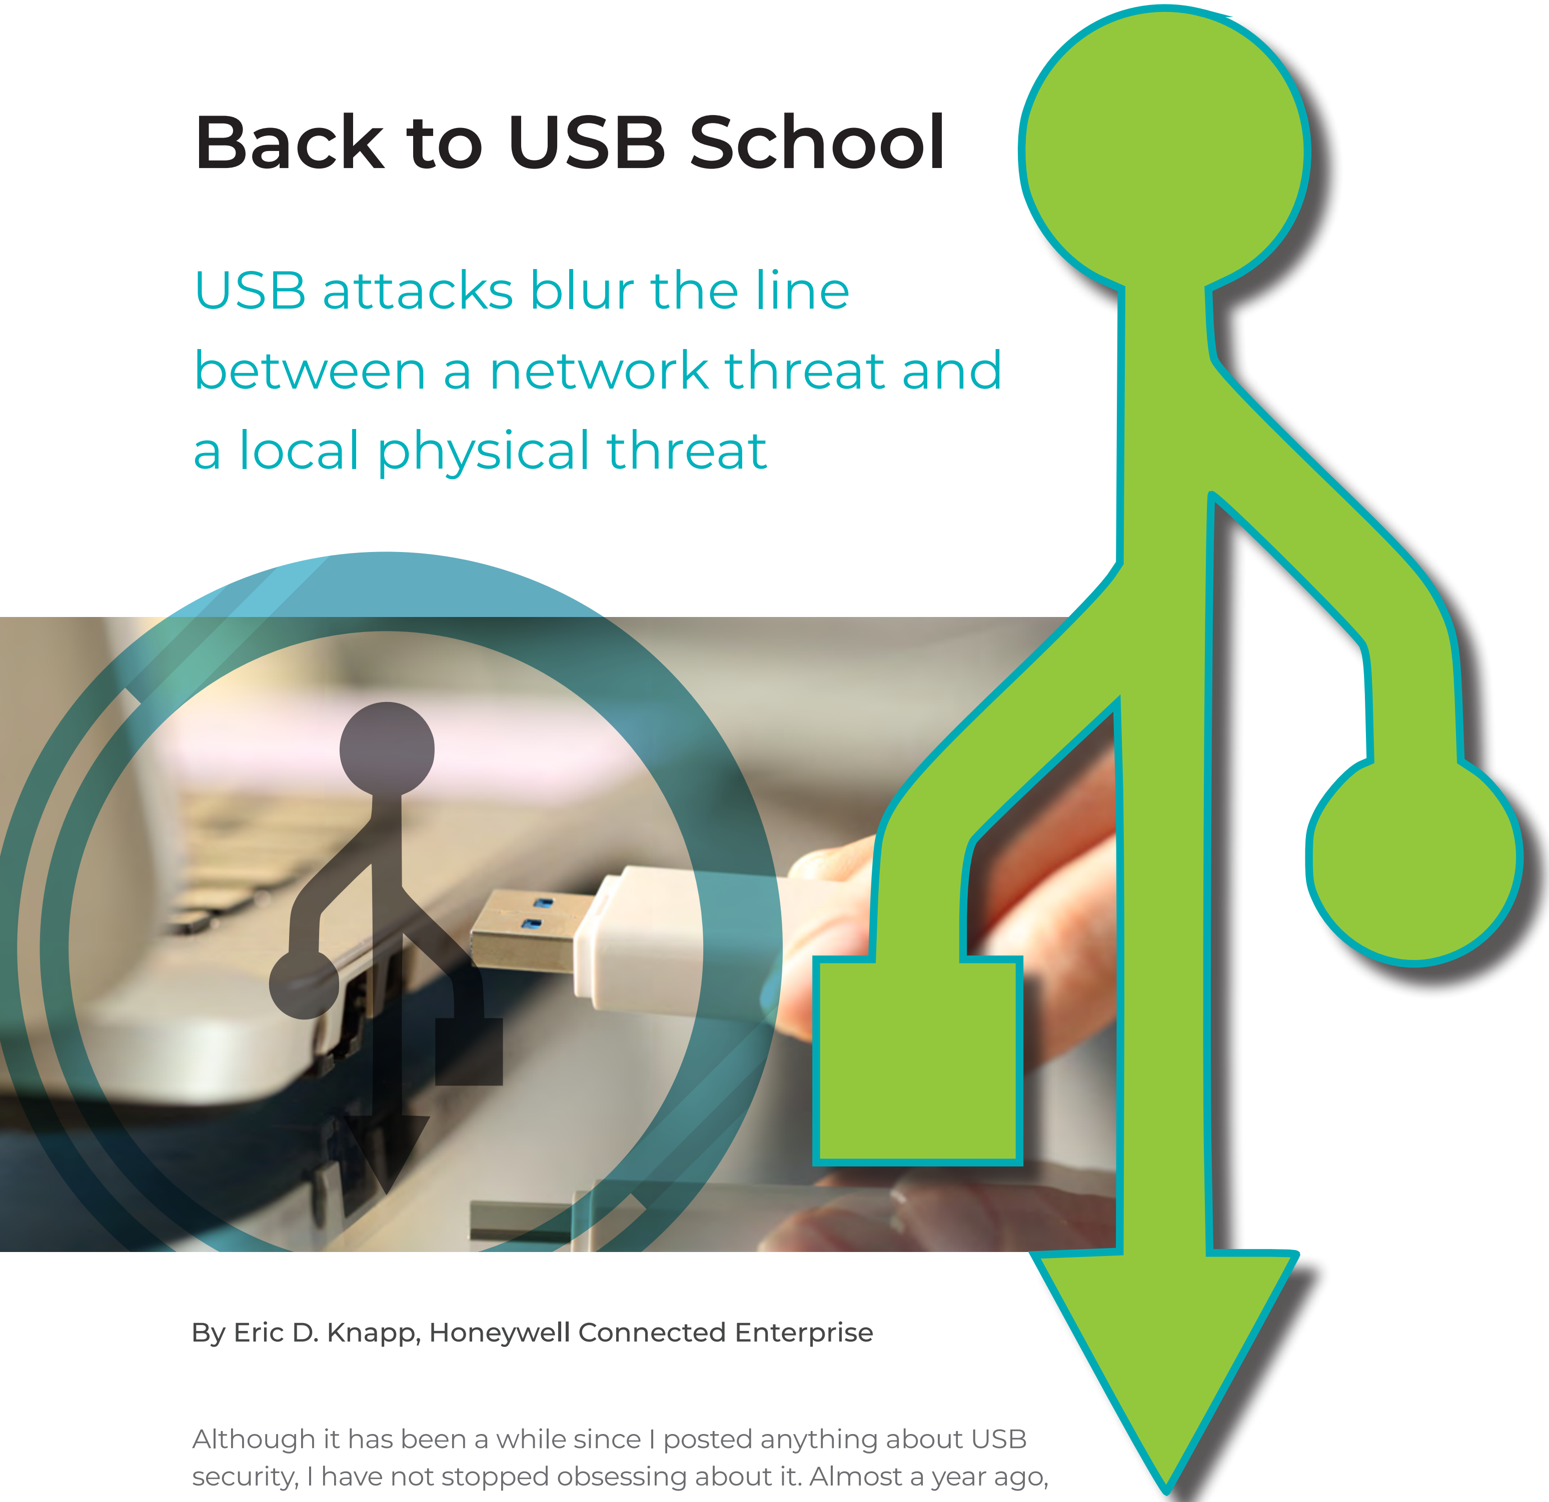

By Eric D. Knapp, Honeywell Connected Enterprise

Although it has been a while since I posted anything about USB security, I have not stopped obsessing about it. Almost a year ago, when the world's biggest health crisis was the voluntary inhalation of vaporized nicotine, I spoke about how easy it was to use USB devices as attack vectors.

There was a small moment during that presentation when I talked about the scope of the USB vector. I mentioned it almost offhandedly, but it has been nagging at me ever since. What I said was, "every USB interface can be home to a whole network of devices, all communicating in a way that is eerily similar to Ethernet. And because almost everything out there today has at least one if not many USB interfaces on it, the USB protocol *essentially* extends every one of our networks into something exponentially bigger."

I probably did not say it that eloquently, but the point was that USB devices and hosts communicate over a network (the "bus" of universal serial bus), and that in turn is connected to every other network. Now vaping has taken back seat to a global pandemic that has everyone doing everything possible remotely . . . and I keep thinking about what this means for the threat landscape, and specifically for USB threats.

So while I felt pretty smart back at RSA Conference, it is time to go back to school—on the bus? Get it? Things have changed a lot in just one year. And, as the security industry has continued to evolve at its typical lightning speed, I have been thinking back to that offhand comment about USB being a network, and how that network could potentially interact with other networks.

> **The USB protocol** essentially extends every one of our networks into something exponentially bigger.

My worry is that everything out there in the world today is connected, and everything out there in the world today has USB. Because a single USB host can connect to dozens of logical interfaces, and each of those can have multiple end points, there could be any number of things on our networks that we do not really know about. It is like every traditional network node is carrying another tiny network around with it. If there was an easy way for an attacker to move freely between the Ethernet network and the USB bus, it would mean there is a new softer, less secure "edge network" coupled to our infrastructure *that we are not even paying any attention to.*

It is a daunting theory, and unfortunately one that in the past year has become a reality. Sure, it was always technically possible, but over the past year it has become not only real but really easy.

At DEF CON 2019, there were two new wireless USB platforms available for purchase (that I am aware of), and at least one more platform was introduced in one of the demo sessions. USB attacks are becoming more interactive, and they are starting to blur that line between a network threat and a local, physical one.

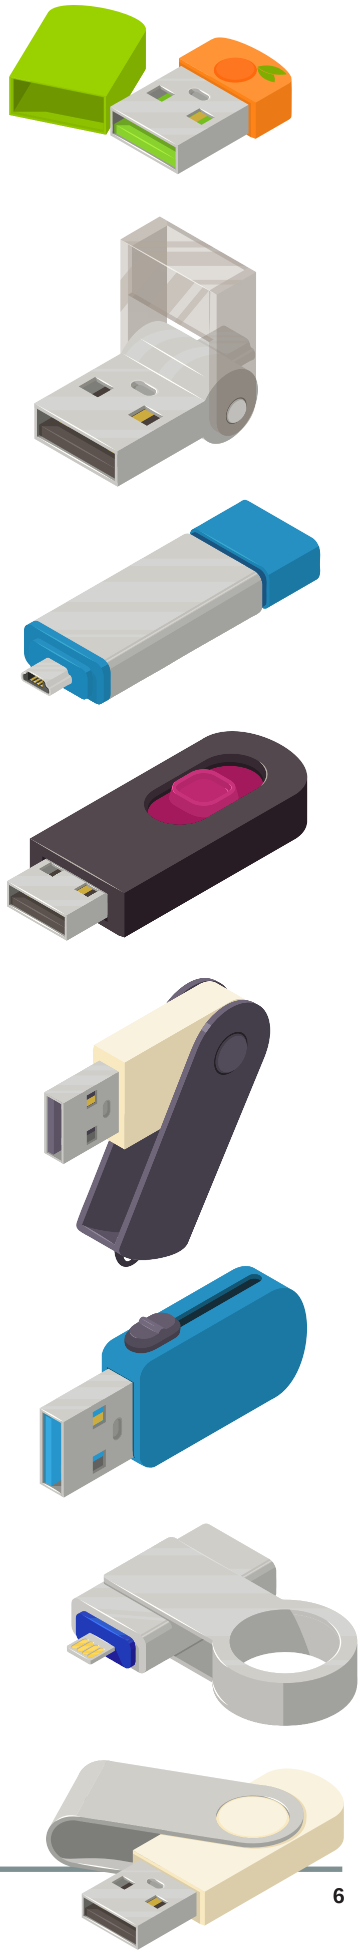

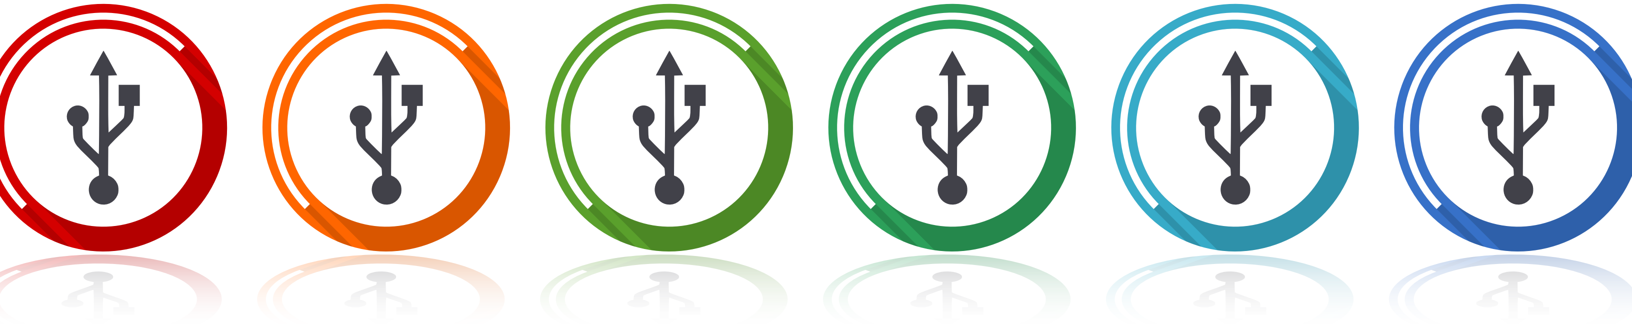

To show just how easy it can be, I wrote this article in notepad remotely, by sending commands over a network to an O.MG cable—a clever and powerful pen-testing tool that hides a tiny server inside a USB cable. That cable was connected to my laptop (as a human interface device, or HID), but also to Wi-Fi. It is a silly example, but one that easily proves that you can remotely influence computers via locally attached USB devices.

What can we do about it? Well, we can and should continue to experiment and learn. To that end, we have been planning on hosting a USB threat challenge later this summer (although that may need to be virtual now, or postponed) to see how clever the hacking community can get.

My personal hope is to see just how far we can push the boundaries using USB as a vector. Armed with that knowledge, we can find new and better ways to cope with this rapidly developing threat vector.

ABOUT THE AUTHOR

Eric D. Knapp (@ericdknapp) is a Senior Fellow at Honeywell Connected Enterprise, where he drives advancements in industrial cybersecurity as the leader of the Global Research, Analysis, and Defense team. Knapp is a recognized expert in industrial control systems cybersecurity. He is the author of *Industrial Network Security: Securing Critical Infrastructure Networks for Smart Grid, SCADA, and Other Industrial Control Systems*, and the co-author of *Applied Cyber Security for Smart Grids*. Knapp has more than 20 years of experience in OT cybersecurity and holds multiple patents in the areas of risk management, asset protection, and secure data transfer. Prior to Honeywell, Knapp held technology leadership positions at NitroSecurity, Intel Security/McAfee, and Wurldtech, and is active on numerous industry boards and committees. Knapp's research and development efforts are the result of his never-ending quest to improve the field of industrial cybersecurity.

Online Learning

Learn more about ways to protect industrial control systems from today's most common cyberattacks, including USB port intrusion and phishing. The latest webinar hosted by Honeywell and Automation.com is entitled "9 Key Ways to Protect Yourself from the #2 Threat to OT Environments" and is available on-demand on Automation.com. Additional materials can be found on Honeywell's BEcybersecure.com website.

tenable.ot™
Powered by Indegy

Get the **Operational Technology Security** You Need. **Reduce the Risk** You Don't.

Complete Visibility

Go deeper than simply listening to network traffic by actively querying devices in their native protocols.

Proactive Risk-Based Insights Into Vulnerabilities

With real-time information you'll always know your risk profile and be ready to address new threats as they emerge.

Unified OT and IT Security

Unify your OT and IT security with a single vendor in an integrated solution.

Gain Full Visibility, Security and Control with Tenable.ot:

Tenable.ot Is the Leader In Industrial Cybersecurity, with Patented Active Query Technology

REQUEST A DEMO

Learn about the 7 Most Common Unsafe Gaps In Industrial Cybersecurity and How You Can Protect Against Them

VIEW INFOGRAPHIC →

The 5 Things You Need to Know About IT/OT Convergence

WATCH WEBINAR →

tenable®

The Increasing Attack Surface: Industrial Environments at Risk

By Michael Rothschild, Tenable

Facing new complex threats to once-separate IT and OT systems, industrial organizations are "de-siloing" to secure their global environment

IT professionals will look at you like you are crazy. Tell them: "The computer running the electrical grid has not been touched in 20 years."

Or tell the IT people in a bottling facility the computer that runs the plant was last moved back during Y2K preparations.

They simply will not believe you.

Why? Because IT and operational technology (OT) approaches to operations are polar opposites. While innovation and security form IT's foundation, OT is more about letting a system run reliably and with as little change as possible. The chasm between IT and OT has traditionally been wide, but it won't be for much longer.

Where did it all begin? IT and OT typically never intersected. IT environments were always connected to the Internet, intranet, and beyond. As a result, security risks were a concern for arguably more than three decades. There was an overarching business and technological need for complete visibility, security, and compliance, mostly because just one attack could shake customer faith, shareholder confidence, and ultimately put a business at risk.

This environment required professionals to constantly evaluate technology and swap it out with a change out schedule generally every 18–36 months—and that was just to stay ahead of the "breach curve."

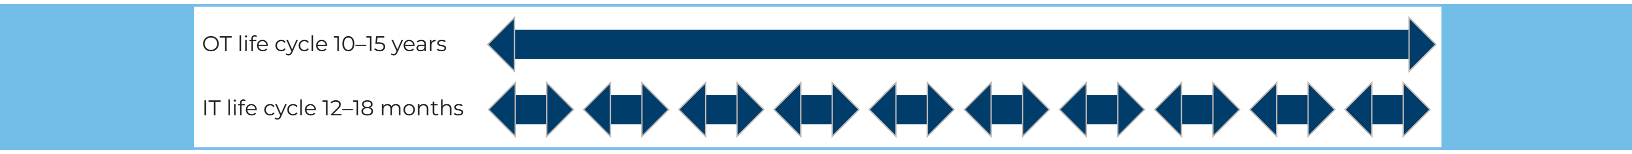

Figure 1. The difference in technology change outs between IT and OT equipment became known as the "life-cycle disparity."

OT environments were a different beast. Whether it was the industrial controller running the cooling tower, purification process, blast furnace, electrical grid, or any number of things, these "old reliable" systems were completely disconnected from everything else.

Security was not a concern, and compliance was not an issue, because "air gaps" separated industrial networks from the rest of the world. Industrial networks were not connected to business networks or the Internet. Because of this "set it and forget it" attitude, technology seldom—if ever—needed switching out. It was not uncommon for OT equipment to be as old as the plant.

A new connected world

In today's connected world, "air gapping" is no longer an operationally feasible solution. Many pundits claim the IT and OT chasm started changing when the notion of the Internet of Things (IoT)—or more appropriately the Industrial Internet of Things (IIoT)—started ramping up.

Industrial and critical infrastructure environments now wholly adopt connecting IT and OT, while also leveraging IoT technology to realize efficiencies and cost savings across the entire organization.

Global connectivity makes our lives easier in many ways. We need not go any further than our smartphone to summon practically anything we want. So too, in the world of critical infrastructure and manufacturing, the ability to have instant access to production lines, manufacturing facilities, and electricity-generating plants enables up-to-the-minute metrics, full visibility, and the ability to control changes from anywhere—with less effort and cost.

Need for anywhere access

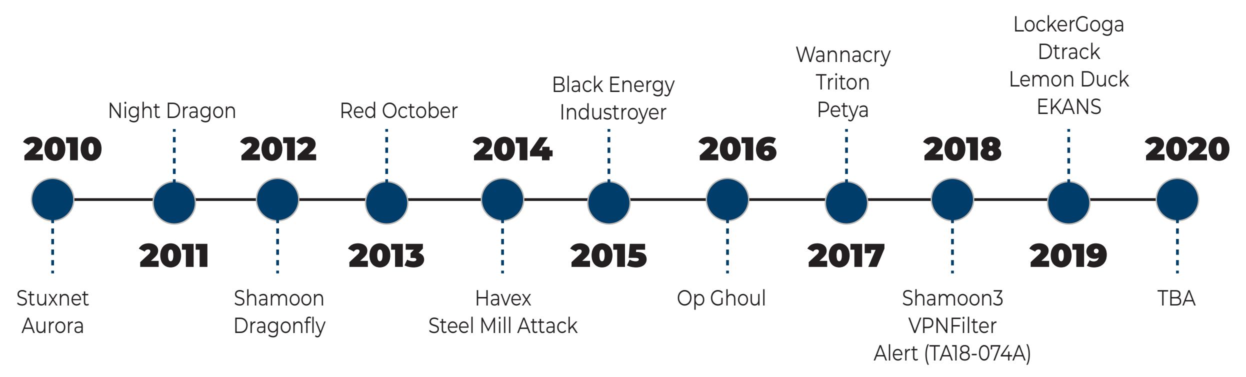

Figure 2. A history of some of the major attacks on OT infrastructures. Courtesy of Tenable

A concrete manufacturing company found it took two days to fire up its blast furnace to the right temperature. What better way of monitoring the progress than over a single pane of glass graphical user interface from anywhere in the world?

Connecting and interconnecting the electrical grid between suppliers was contributory in causing great blackout of 2003 where almost the entire Northeast was without power due to a cascading failure. Systems could not talk to each other and the visibility needed to thwart this failure was simply not there.

In each of these cases, businesses found practical applications to connect things to the web. The results were huge benefits from cost saving, visibility, and efficiency perspectives.

These cases heralded the convergence of once separate IT and OT systems, as well as rapid adoption of IoT technology by industrial organizations. But convergence created a new problem—connected industrial controllers had little in the way of defense against cyberattacks. Little did people know that they would quickly find OT devices in the crosshairs of attacks that could alter the way we live.

Cyber convergence challenges

The increase in the number of cyber incidents on industrial control system (ICS) networks is a reality we can no longer ignore. Few argue against the attack surface changing to encompass both IT and OT. Because these two different worlds are now connected, an attack that starts on an IT environment can quickly move to an OT environment and vice versa.

Lateral movement is a preferred attack methodology for hackers. It is relatively easy to find a weak link in the system, leverage it as a point of entry, and then quickly own the entire network. Much like when one person sneezes in a room, a bunch of people can get a cold, so too our interconnected systems across IT/OT environments share “cyber germs” that can take down an entire system.

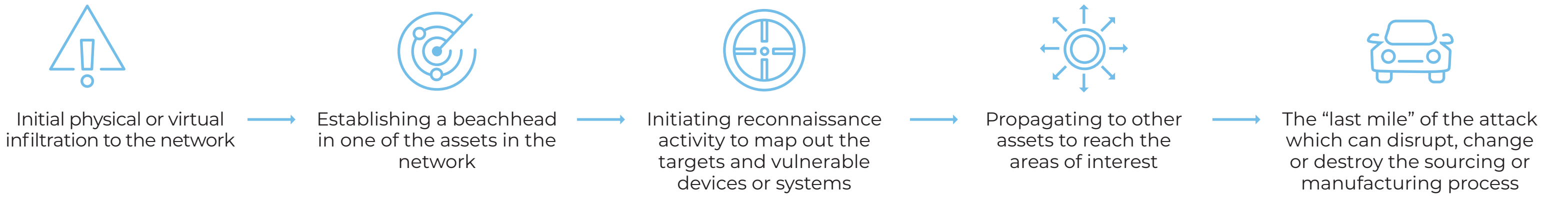
Initial physical or virtual infiltration to the network
Establishing a beachhead in one of the assets in the network
Initiating reconnaissance activity to map out the targets and vulnerable devices or systems
Propagating to other assets to reach the areas of interest
The “last mile” of the attack which can disrupt, change or destroy the sourcing or manufacturing process

Figure 3. Typical cross-platform attack etiology

Protection against ICS threats

Bringing these two substantially different IT and OT worlds together is a challenge. To address new complex threats that broaden attack surfaces and increase the amount of attack vectors, organizations are “de-siloing” their approach to securing their global environment.

What is necessary to help address clear and present needs?

1. **360-degree visibility:** Attacks can easily propagate in an IT/OT infrastructure. With a single platform to manage and measure cyberrisk across OT and IT systems, you will gain complete visibility into the converged attack surface. You will also want a solution that natively integrates with leading IT security and operational tools, such as a security information and event management (SIEM) solution, log management tools, next-generation firewalls, and ticketing systems. Together, this builds an ecosystem of trust where all your security products can work together as one to keep your environment secure.
2. **Threat detection and mitigation:** Ensure your OT security solution leverages a multi-detection engine to find high-risk events and behaviors that can impact OT operations. These engines should include the following detection capabilities:
    - *Policy-based:* Activate predefined policies or create custom policies that whitelist and/or blacklist specific granular activities that may indicate cyberthreats or operational mistakes that trigger alerts. Policies can also trigger active checks for predefined situations. This is crucial for discovery of risky events that do not rise above the statistical noise (e.g., malware, reconnaissance activity, querying device firmware versions from a human-machine interface [HMI]).
    - *Behavioral anomalies:* Where the system detects deviations from a network traffic baseline based on traffic patterns. Pattern baselines include a mixture of time ranges,

protocols, devices, etc. Among other things, it allows detection of suspicious scans indicative of malware or rogue devices in your network. It will also help detect Zero-day attacks where no policy or signature has yet been created.

- *Signature updates:* By leveraging a crowdsourced signature database (such as Suricata), you can detect attacks throughout all stages and get alerts with context about suspicious traffic that can indicate reconnaissance, exploits, installed malware, lateral propagation, and more. The threat detection engine should ingest new signature updates to address new threats as they evolve.

Industrial and critical infrastructure environments are increasingly converging IT and OT, while also leveraging IoT technology to realize efficiencies.

3. **Asset inventory and active detection:** Your OT security solution should provide unparalleled visibility into your infrastructure—at the network level down to the device level. It should combine native communication protocols to actively query IT, as well as OT devices in your ICS environment, to identify all activities and actions, as well as gaining deep situational awareness across your network and devices in your network.
4. **Risk-based vulnerability management:** Drawing on comprehensive and detailed IT and OT asset tracking capabilities, your OT security solution should generate a prioritized list of vulnerabilities and risk levels for each asset in your ICS network. These reports should include risk scoring and detailed insights, along with mitigation suggestions. Vulnerability assessments should include parameters, such as firmware versions, relevant CVEs, proprietary research, default passwords, open ports, installed hotfixes, and more.
5. **Configuration control:** Should track any changes made to OT assets, whether they are user executed or malware based via the network or by local connection. This capability should provide a full history of device configuration changes over time, including granularity of specific ladder logic segments, diagnostic buffers, and tag tables. This enables administrators to establish a backup snapshot with the "last known good state" for faster recovery and compliance with industry regulations.

Collaboration between IT and OT can help mitigate risks and vulnerabilities that traverse these two unique and now deeply intertwined infrastructures. By combining a strong IT security posture with an equally strong OT posture, industrial organizations can protect ICS networks from external and internal cyberthreats—now and in the future.

ABOUT THE AUTHOR

Michael Rothschild, with more than 20 years of security experience, is the senior director of OT Solutions at Tenable. He is a past professor of marketing and has published a number of works on the topic. He currently occupies an advisory board seat at Rutgers University. With a passion for healthcare, Rothschild is a volunteer EMT and deputy chief of his local ambulance corps. He also teaches for the American Heart Association.

Ready for the Next Big Step in EtherNet/IP?

Get to know CIP Security, a secure standard for the transport of EtherNet/IP messages over an industrial network

By John S. Rinaldi,
Real Time Automation

Encryption dates as far back as the Spartans of ancient Greece and possibly even further. Like every army before and after them, the ancient Spartans needed a mechanism to send confidential messages to field commanders. Their solution was to secure their communications by wrapping a long piece of leather around a wooden rod and writing a message vertically on the leather. When unwrapped from the wooden rod, you had a piece of leather inscribed with a series of seemingly senseless letters. The leather could then be carried to the intended recipient who, knowing the diameter of the rod (the "key"), would wrap the leather around another rod of the same diameter and read the message. If it fell into enemy hands, anyone lacking the "key" would have a nonsensical series of letters. Of course, this was far from foolproof; brute force decryption—trying wooden rods of different diameters—could eventually decode the message, but it was ingenious for that era.

While we are a long way from wrapping wooden rods with strips of leather, the need for confidentially exchanging messages has not changed. Today, we have Ethernet systems on our factory floor exchanging messages between controllers and end devices. In the past few years, these Ethernet systems have been extended to link enterprise and cloud applications to the factory floor.

Unfortunately, extending connectivity beyond the factory floor has increased the vulnerability of those systems to cyberattacks. Attackers, sensing an opportunity, have shifted their attention from personal computers and servers to the world of factory automation. Because the majority of these attacks are not publicized, no one knows for certain how many plants have had their servers locked, important data stolen, messages altered, and programmable controllers hijacked.

In the past, it was not uncommon to have insecure controllers directly connected to the Internet. Over the years, these controllers have been removed, updated, or replaced with newer versions that are more cybersecure. Most manufacturing installations have also added defense-in-depth strategies that make it much more difficult to get to controllers and I/O networks from the outside. What is often still open and vulnerable, though—if you can get to it—is the inside, the I/O network side of programmable controllers.

If you can get access to an EtherNet/IP, Modbus TCP, or PROFINET IO network, you can often have free reign to create all kinds of havoc. There is generally nothing stopping you from accessing the controller tags over that network: turning pumps on or off, increasing motor speeds, or opening and closing valves.

Even with strong cybersecurity protection from the outside, factory floor systems can be compromised from the inside. Most facilities have an army of Internet of Things (IoT) vendors, automation vendors, technicians, system integrators, and corporate engineers who come on site and knowingly or unknowingly bring viruses, malware, time bombs, and worse into your plant and onto your critical I/O networks.

If you can get access to an EtherNet/IP, Modbus TCP, or PROFINET IO network, you can often create all kinds of havoc.

EtherNet/IP, the Ethernet implementation of the Common Industrial Protocol (CIP), was never designed as a secure communications transport. It is designed for ease of use and flexibility. Anyone can make connections to an EtherNet/IP adapter and execute any operation, including a reset of the device. This makes EtherNet/IP a very insecure communications protocol.

In light of this, ODVA recently began deployment of CIP Security for EtherNet/IP. CIP Security is a secure standard for the transportation of EtherNet/IP messages. It allows communication between trusted entities, and disallows communication between untrusted entities on an EtherNet/IP network.

This article introduces CIP Security for EtherNet/IP. It explains what is meant by CIP Security and describes the technologies that it is based on, the new CIP objects that are required, and how developers and end users should move forward in this era of secure EtherNet/IP.

What is CIP Security?

CIP Security is designed to protect not only EtherNet/IP adapter devices (end devices) from access by unauthorized parties, but also to protect programmable controllers. Attackers have noted that programmable controllers are more resilient to outside entities (Internet attacks) than to inside entities (I/O network attacks).

The ODVA designed CIP Security to protect programmable controllers and devices on I/O networks from attacks originating on those networks. At first blush, attacks on the I/O network seem unlikely. I/O networks are not generally connected to the Internet, so what is the concern? In practice, these kinds of attacks are not all that unlikely.

Contractors come and go from a facility and connect to networks with laptops that may be compromised. Employees may fail to disable open ports on switches. Some employees knowingly engage in sabotage. There is a myriad of ways for attackers to get access to your I/O network. CIP Security is designed to increase the immunity to such attacks. The secure EtherNet/IP transport provides the following security attributes:

- **Authentication of the end points:** ensuring that the target and originator are both trusted entities. End point authentication is accomplished using X.509 certificates or preshared keys.
- **Message integrity and authentication:** ensuring that the message was sent by the trusted end point and was not modified in transit. Message integrity and authentication is accomplished via TLS message authentication code (HMAC).

- **Message encryption:** optional capability to encrypt the communications, provided by the encryption algorithm that is negotiated via the TLS handshake.

A fundamental design tenet is that not all devices on an EtherNet/IP network need the same level of protection. Some devices are less critical, and some are more critical to an automation system. The required protection is not identical; CIP Security defines two security profiles to offer that different level of protection.

EtherNet/IP Confidentiality Profile provides secure communications by requiring authentication and data integrity for all EtherNet/IP messages. Authentication means that an EtherNet/IP device identity is verified to be the device it claims to be. Data integrity means that the data within the EtherNet/IP message can be relied upon to be accurate and consistent.

EtherNet/IP Authorization Profile goes one step further than the Confidentiality Profile. It provides user authorization. With the authorization profile, an application requesting an action like opening or closing a valve would have to be authorized to take that action.

EtherNet/IP Devices that do not support CIP Security can coexist with devices that support the Confidentiality or Authorization Profiles.

> **CIP Security** is designed to protect not only EtherNet/IP adapter devices (end devices) from access by unauthorized parties, but also to protect programmable controllers.

The two CIP Security trust models

A trust model is a very important consideration in manufacturing system security. The trust model is the collection of rules that govern how a device decides to trust another device. A trust model that is too soft (flexible) cannot provide the integrity, confidentiality, and authenticity that you need. A trust model that is too hard (inflexible) becomes such a burden on daily operations that it lowers your productivity. But a trust model that is “just right” provides you with that integrity, confidentiality, and authenticity without impeding legitimate entities that need to communicate and keep product flowing out the door.

CIP Security supports two trust models: preshared key (PSK) and certificates. Both are useful. Both have advantages and disadvantages.

PRESHARED KEY (PSK): Preshared key is an uncomplicated system that works well in small systems. Private key sharing operates very simply. A private key is known and shared by all the devices in a network. The key is used to encrypt messages. Any device that knows the private key is authenticated and able to encrypt and decrypt messages. For added protection, the key is changed at some set interval, sometimes as part of a maintenance cycle.

X.509 CERTIFICATES: X.509 certificates are a standard way for two devices to securely communicate. Each device has a certificate identifying the entity issuing the certificate. That entity can be the device itself (self-signed certificate), the vendor who manufactures the device, or some outside authority that is trusted by other devices with which it wants to communicate. The device receiving the certificate can send encrypted messages to the originator by encrypting the message with the public key in the certificate. The private key, which is never disclosed outside of the device, is used to decrypt the message encoded with the public key.

CIP Security vendors are required to support both trust models. End users can decide which makes more sense for their facility and commission their device appropriately.

How is CIP Security applied to EtherNet/IP?

EtherNet/IP, like PROFINET IO and other industrial protocols, uses both acyclic and cyclic communications. Acyclic communications are used for moving information between scanners (controller-side devices) and adapter devices (I/O-type devices). Cyclic communications are used for moving I/O data between scanners and adapters. Acyclic messages send configuration and information like ramp-up time on an intermittent schedule, while cyclic communications are repeated on a continuous basis.

The underlying communication layers are different for both communication types. Acyclic communication uses TCP messaging to move messages. TCP communication is reliable. Packets are sent (and received) in sequence, and the sender gets an acknowledgement that each packet arrived on time. Cyclic communication uses User Datagram Protocol (UDP). UDP is a fire-and-forget protocol without any confirmation that the packets arrived at all and in what order.

Because of the underlying differences of these transport layers, CIP Security uses two different security mechanisms. Acyclic communications (TCP transport) are secured using Transport Layer Security (TLS), while cyclic communications are secured using Datagram Transport Layer Security (DTLS).

TLS is a well-known Internet security standard designed to ensure message integrity, to authenticate end points, and to keep the contents of messages private. You are using TLS whenever you see that little lock and the "https:" at the beginning of a URL.

The nature of cyclic messaging and UDP makes TLS unsuitable for EtherNet/IP cyclic communication. DTLS is a variant of TLS that is designed for cyclic communications. It implements:

- a retransmission timer for lost messages
- message sequence numbers to queue messages properly
- fragmentation of large messages
- replay detection to discard previously processed messages

EtherNet/IP cyclic messaging relies on DTLS to ensure privacy, authentication, and data integrity for the I/O messages that are so important to a properly functioning automation system.

The three new EtherNet/IP objects

All the Common Industrial Protocol technologies—EtherNet/IP, CompoNet, ControlNet, and DeviceNet—are object-based technologies. That means that users interact with CIP devices by interacting with the objects implemented in those devices. Three new objects provide the object model access to CIP Security for EtherNet/IP:

- **The CIP Security Object description:** The CIP Security Object is a high-level control object. It is the simplest of the CIP Security objects. It provides a flag that external entities can check to determine if a device is in the CIP Security configuration state. Instances: The CIP Security Object supports a single instance.

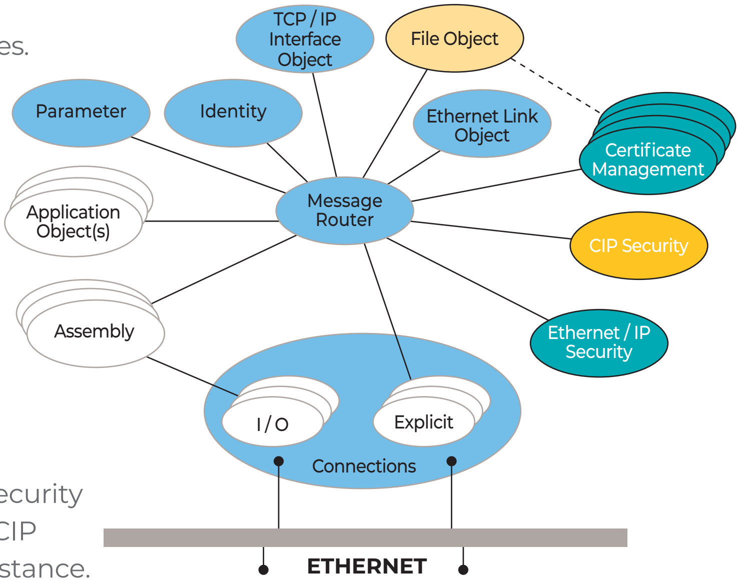

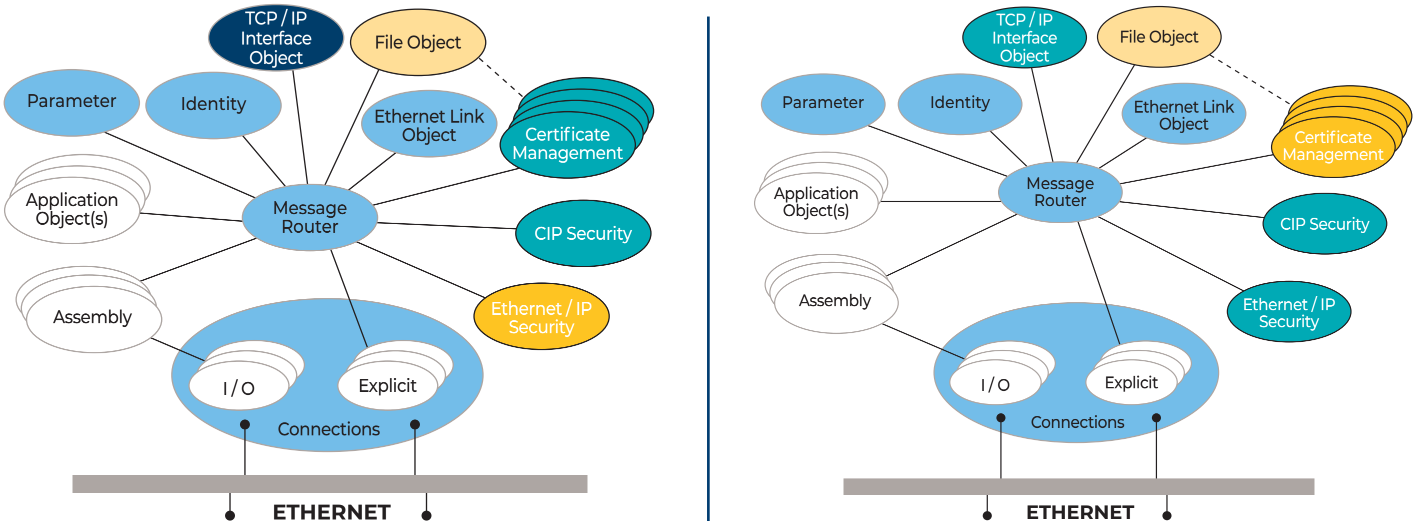

- **The EtherNet/IP Security Object description:** The EtherNet/IP Security Object is the CIP Security object that manages the parameters that govern how CIP Security operates on an EtherNet/IP device. It manages the parameters that control TLS and DTLS operation, the cipher security suites, the lists of trusted authorities, and the mechanisms for obtaining X.509 certificates. It identifies the current active device certificate that the device is using for secure communications. Instances: The EtherNet/IP Security Object supports a single instance per TCP/IP object.
- **The Certificate Management Object description:** The Certificate Management Object (CMO) is the CIP Security object that manages the X.509 certificates maintained by the device, and creates Certificate Signing Requests (CSR). Signing requests are applications to a Certificate Authority for creation of an X.509 certificate. In some commissioning applications, a configuration tool will request the CMO to create a signing request. The CMO stores the request in the file object where it can be read by the configuration tool and used to obtain a certificate from a Certificate Authority local to the application. Instances: Unlike the other CIP Security objects, there are multiple instances of the CMO. Instance 1 manages the default X.509 certificate, while additional instances manage any additional certificates loaded into the device.

What's next? Now that Rockwell ControlLogix supports this security mechanism and is rolling out products, CIP Security over EtherNet/IP may soon become a checklist on the purchase specification for all EtherNet/IP adapter devices. This means everyone with an EtherNet/IP device is going to need to upgrade their EtherNet/IP adapter to support secure transport.

ABOUT THE AUTHOR

John S. Rinaldi is chief strategist and director of WOW! for Real Time Automation (RTA) in Pewaukee, Wisconsin. Contact him at http://www.rtautomation.com/contact-us or https://www.linkedin.com/in/johnsrinaldi. With a focus on simplicity, support, expert consulting, and tailoring for specific customer applications, RTA is meeting customer needs in applications worldwide. Rinaldi is not only a recognized expert in industrial networks and an automation strategist, but also a speaker, blogger, and the author of six books and more than 100 articles on industrial networking. To get EtherNet/IP with the CIP Security add-on, contact Real Time Automation at 1-800-249-1612 or sales@rtautomation.com. For more information on EtherNet/IP, a good place to begin is the book, *EtherNet/IP: The Everyman's Guide to EtherNet/IP*. You can get it at no charge by visiting https://www.rtautomation.com/rtafreegift. Use "113" as the publication code.

Bridging the IT and OT Cybersecurity Divide

By Peter Vescuso,
dragos.com

Experts from both domains can bolster business resiliency no matter what the cyber-threats target

Industrial organizations and modern enterprises are grappling with a two-sided cybersecurity problem. They must learn to take a mature security posture in both their information technology (IT) and operational technology (OT) environments at a time when both are coming under increasing attacks—and as the line between the two realms blurs together more and more by the day.

The challenge is that while OT shares some similar operating systems, network connections, digital architectures, and cybersecurity risks as IT, there is definitely not a one-to-one relationship between the two worlds. There remain many unique constraints to securing the operational world of industrial control systems (ICSs), which means that organizations cannot simply copy and paste IT cybersecurity strategy for OT cybersecurity.

Nevertheless, IT and OT networks are increasingly interconnected to support digital transformation efforts and initiatives that drive Industry 4.0, which means accountability and priorities need to be unified. Plus, organizations can still learn a lot from the long evolution of IT cybersecurity threats and defense. Applying those lessons to OT and tailoring that knowledge to the operational environment can help create an OT cybersecurity strategy that meets the threats and circumstances of ICS security both today and in the future.

However, that can only be done if organizations open the lines of communication between IT and OT. Experts from both domains must start to work cohesively to bolster the resiliency of the business no matter which side of the house the cyberthreats target.

Why OT cybersecurity matters

Digital transformation. Enterprises are spending trillions on digital transformation today, and industrial applications are at the spear tip of these investments. When industrial concerns use cloud-connected software to better automate plants, bolster predictive maintenance, or connect industrial devices at the edge to business intelligence platforms, they are by definition more tightly coupling OT with IT systems. The business benefits are tremendous, but the process of digitally transforming industry also greatly expands the cyberrisk to the OT environment.

The world is industrial. Although the field of industrial systems has never been just about power plants and manufacturing facilities, even the perception of that no longer exists. Whether it is OT systems that track shipping operations, smart heating and lighting systems that run office complexes, smart robots that stock store shelves, or automation systems that streamline warehouses, operational technology is everywhere in the enterprise today.

These are the systems that make up the fabric of our real-life business worlds—ones that would put business continuity or people's safety at risk if they were compromised. And yet they are often forgotten from a cybersecurity perspective.

Attackers are already here. One of the biggest problems enterprises face in bridging the IT to OT cybersecurity divide is complacency. There is a perception that because the industry has not yet witnessed evidence of cyberattacks in the OT environment, it must not need OT cybersecurity monitoring. The common mantra is "There's no way our OT is a target—we have not seen any attacks."

The thing is that many attackers operate stealthily, and enterprises just do not have the mechanisms in place to see them within their OT systems. This breeds a scenario where organizations lack cyber-visibility. Because they do not monitor OT, to them the adversaries do not exist. However, time and again, Dragos runs assessments for new customers that uncover adversaries who have been present in the OT environment all along.

The OT cyberthreats of today and tomorrow

OT cyberthreats are both worse than you realize and not as bad as you want to imagine.

Without a doubt, enterprises must take ICS and OT security seriously, because the compromises are quietly accelerating. Publicized examples of successful attacks against OT systems remain remarkably rare, because most in the OT cybersecurity community understand that it is better for the ICS world and public safety to keep successful attacks under the radar. Within individual organizations, many stakeholders may be unaware of a problem, because when accidents or maintenance events with cybercomponents strike, they are often undiagnosed as cybersecurity incidents.

But these incidents and the perpetrators who carry them out are growing more prevalent. In this regard, the OT threat environment mirrors its IT threat cousins. Over the decades, IT threats have grown more prolific and more sophisticated. A similar evolution is slowly unfolding within OT. Whereas a few years ago we would see maybe only one or two global adversary teams capable of carrying out attacks against ICS systems, Dragos now tracks 11 groups that are persistently targeting OT assets around the world. And there are more threat actors and capabilities brewing.

At the same time, the larger cybersecurity community and the early advocates for OT cybersecurity must slow down the hysteria. The claims that phishing emails will take down power grids are overwrought and hurting the cause. First of all, the ICS community on the whole has built out a very resilient physical infrastructure. The beauty of those global efforts by engineering and operations professionals to advance industrial safety is that this focus has already led to a natural level of security within so many OT systems.

Additionally, the saving grace for the cybersecurity of OT systems today is that most of them are still very custom and very heterogenous. True, many OT systems run Windows like their IT cousins. But in OT there still exist many customized processes, customized hardware, customized embedded systems. Just by this very design it takes attackers a lot more effort, a lot more reconnaissance, and a lot more data collection to figure out how to build malicious software to achieve their attack goals. Most importantly, it blocks attackers from scaling attacks, because they cannot easily port techniques from one facility or organization to another.

The point is: Do not panic—but be aware that the mitigating factors for OT cybersecurity will start to deteriorate in the coming years. As digital transformation accelerates, industrial control systems will grow more homogenous, more connected, and more converged with IT. For example, cloud convergence has many organizations moving toward cloud-direct connections to historians and sensors. This opens up the kind of back doors into the OT environment that no one is properly planning for or thinking through.

As OT infrastructure changes through digital transformation, the threat actors will adapt to that with greater sophistication. Thus, it becomes crucial to add a higher level of cybersecurity competency and controls to the mix of safety measures already present in the industrial environment.

What we can learn from IT cybersecurity

As OT cybersecurity threats begin to advance, organizations can certainly learn to defend against them by looking at how IT attacks and defensive philosophies have evolved over the years. In the past decade, the IT networks have been increasingly deluged with automated attacks on all sides, perpetrated by adversaries with numerous and complex motivations. In an era of rampant ransomware attacks, financially motivated attackers are carrying out cyberespionage, theft, disruption, and destruction of IT assets.

The best practitioners in IT cybersecurity have recognized that this constant and persistent attack pressure means that it is inevitable that the bad guys will eventually manage to break into the network—somewhere, somehow. But the best cyberdefenders came to the dual realization that this does not have to translate to adversary success in achieving their attack objectives.

IT security veterans know that the goal is not to keep threat actors from ever exploiting vulnerabilities in any given system. It is to keep them from stealing valuable intellectual property, committing fraud, encrypting machines for ransom, and so on.

The fundamental truth in IT cybersecurity today is that the most resilient cyberdefenses are those that slow down adversary progress in the network and that speed up incident response to the initial break-in. It has become survival of the fastest, and veterans in IT cybersecurity have found that digital resilience boils down to three important metrics: time to detect, time to investigate, and time to remediate.

These metrics are in direct opposition to a concept and attack measurement the IT industry calls “breakout time.” Breakout time is the length of time it takes for an adversary to use an initial foothold on the network to break out of that first system and start attacking other systems in the network.

To counter that, the best in IT cybersecurity strive for the 1-10-60 benchmark. That benchmark dictates that if you can detect attacks in one minute, investigate in 10, and remediate in 60 minutes, you can generally thwart adversaries from ever getting close to their attack objectives.

Now, even in IT cybersecurity, that response speed is a reach goal at best. Most detection, investigation, and remediation response times are measured in hours and days rather than minutes. However, the closer organizations move their metrics toward the benchmark, the more they move the needle on cyberresilience.

The differences between IT and OT cybersecurity

Let’s be realistic. OT cybersecurity is nowhere close to achieving the detection, investigation, and remediation times of the IT world. And that is OK for now.

We should bring the fundamental truth about IT cybersecurity to bear on OT while keeping in mind that OT is very different. In the most simplistic way, you can think of it this way:

OT = IT + PHYSICS

Physics in this equation stands in for the physical processes that OT systems control—whether it is machines and robots in manufacturing facilities, pumps and valves at water stations, or electrical grid equipment run by the power plant.

The physics piece is the hardest part for attackers to influence. It takes quite a bit of planning and design for them to execute manipulations against physical processes and make an impact on facilities and equipment. Take for example the public attack in Saudi Arabia in 2017 using a piece of OT-focused malware called TRISIS. In that example, the adversary had compromised environments for three years before carrying out an attack against an oil and gas facility. This was the first publicly disclosed OT cyberattack clearly designed to injure or kill someone. Fortunately, in this case, the attack failed to hurt anyone due to an error in the malicious code.

However, it does offer a good lens into the problem—namely that there is a magic window for cyberresponse, and it is likely to shrink due to digital transformation and convergence.

At the same time, it is crucial to remember that OT has a different mission, different systems, different threats, and different impact on organizations than IT. Safety, environmental impact, process availability, and intellectual property are key for OT.

Many of the basics of IT security simply do not apply. For example, vulnerability and patch management are fundamental to IT security, but much less important for OT, because many of the vulnerabilities in OT do not necessarily threaten the ultimate safety or mission of that OT system. A recent Dragos study found that some 64 percent of all industrial vulnerabilities do not actually introduce any risk, and a further 34 percent were inaccurate. This means that in the industrial world a patch-at-all-costs mindset does not make sense so much as one that has organizations smartly patching but prioritizing architecture and threat tactics instead.

The overarching lesson is that there are definitely lessons to learn from IT cybersecurity, but as organizations seek to improve OT cybersecurity capabilities it does not make sense to copy and paste your enterprise cybersecurity strategy into the ICS.

Where to get started

Applying lessons from IT cybersecurity and tailoring them to the OT environment is a years-long process toward maturation. But there are some important first steps that organizations can take to kick start their OT cybersecurity strategy and execution.

1. Engage operations

Cybersecurity professionals who want to help improve OT risk postures should start first by listening and learning the language of operations. This can be initiated with a gesture as simple as bringing a box of donuts to break the ice and start a friendly conversation with operators and engineers. Use that opening to ask them to teach you about what goes on in their side of the house. This should be done with no security ulterior motives: no checklists, no enforcement efforts, no vulnerability benchmarks. This opens up a conduit for future cooperation to create relevant cybersecurity policies and procedures that align with OT objectives.

2. Initiate knowledge transfer

The cybersecurity skills gap experienced in the IT world is magnified in OT. It is hard to get access to industrial environments for training purposes, and industrial cyberranges are often extremely costly with few virtualizations. Organizations should be seeking out ways to transfer knowledge and share it—to make more experts in-house and develop security champions among operators and engineers. A good way to initiate that knowledge transfer is to bring in external teams such as Dragos' professional services to do assessments of the environments. Do not just get a report from them—ride along during the assessment and ask lots of questions.

3. Read up and train

Beating the OT cybersecurity skills shortage and learning the language of OT cybersecurity will require all stakeholders to read up and train along the way. Fortunately, the resources are growing for OT cyberdefenders, many of them free. We list a few at the end of this article.

4. Make OT threats visible

The only way to understand the depth and breadth of your OT risk is to start adding better visibility to the OT environment. Use security monitoring to put the right information at the fingertips of defenders, operators, and engineers. But learn from the flubs of IT security in the past—do not overload defenders with every piece of possible information. Be sure systems offer up vetted, relevant, and actionable OT security information so that teams are not drowned out. Bubble up visibility—put information at their fingertips but vet information and make it relevant and actionable—without drowning small teams out.

5. Go on a hunt

Once you have observed, learned the language of OT, grown to love your operations, and learned more about your environment, go on an OT threat hunt. Be proactive in your own environment, and you will start to figure out what you have and what you do not have in terms of information collection and defenses. It is a great way to learn more about the environment and continually improve your risk posture.

RESOURCES

Robert M. Lee's reading list
https://www.robertmlee.org/a-collection-of-resources-for-getting-started-in-icsscada-cybersecurity

Dragos platform
https://dragos.com/platform

Industry news
https://dragos.com/blog/industry-news/a-dragos-industrial-control-system-security-reading-list

SANS ICS courses
https://ics.sans.org/training/courses

Dragos five-day course
https://dragos.com/training

ABOUT THE AUTHOR

Peter Vescuso leads the marketing team at Dragos, which specializes in helping to defend industrial organizations that provide running water, functioning electricity, and safe industrial working environments. This article is a distillation of a recent webinar titled "IT and OT: A bridge too far? Crowdstrike and Dragos don't think so." Vescuso is a seasoned B2B marketing and enterprise software veteran who has spent more than two decades leading marketing for high-growth software businesses. Prior to joining Dragos, he was division vice president at PTC, a $1.5B global software company where he was responsible for marketing digital transformation solutions to the manufacturing industry, including the market-leading industrial IoT platform ThingWorx. Vescuso holds a bachelor's degree in mechanical engineering from New York Institute of Technology and a master's degree in operations research and management from the Thayer School at Dartmouth College.

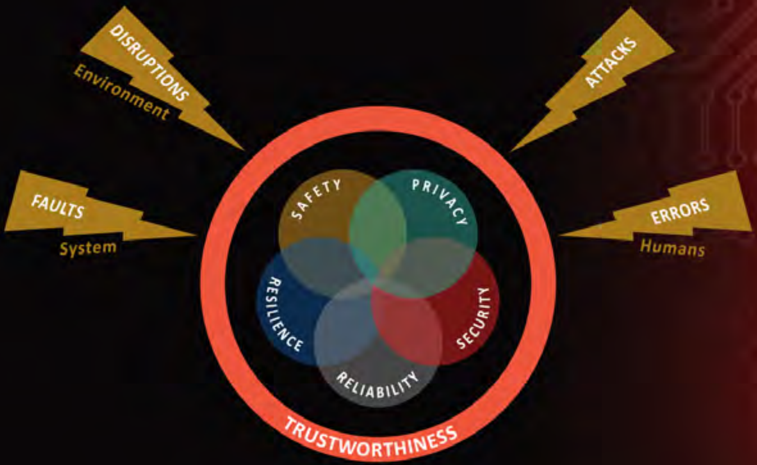
DISRUPTIONS
Environment
ATTACKS
FAULTS
System
ERRORS
Humans
SAFETY
PRIVACY
RESILIENCE
SECURITY
RELIABILITY
TRUSTWORTHINESS

ABOUT THE AUTHOR

Courtney Schneider is cyber-policy research manager for Waterfall Security Solutions, a global industrial cybersecurity company, protecting critical industrial networks since 2007. This article first appeared as a blog post of the Industrial Internet Consortium.

Security at the Edge with Microsegmentation

Industrial and Industrial Internet of Things (IIoT) networks almost always represent engineering risks, as well as conventional "business" risks. IIoT is the ultimate mind meld of information technology (IT) and operational technology (OT) networks. The IIoT connects edge devices in OT networks directly to the Internet to enhance operational efficiencies. What confuses security designs for IIoT deployments is differing kinds of risk.

When IIoT deployments present unacceptable physical consequences, we need strong protections for the edge devices.

OT practitioners and engineers plot risk on a spectrum from unacceptable physical consequences to safe, correct, continuous, and efficient physical operations. Conventional security practitioners, however, focus on protecting information, cyberresilience, incident response, data recovery, and business continuity. Conventional cyberassets are part of a sea of networks, some needing more protection than others, managed for business risk.

What then of IIoT security, which basically melds these two concepts of physical and business risk together: the ubiquity of IT networks layered on physical control and industrial networks? How do we implement a security program to simultaneously satisfy these very different needs from IT, OT, and engineering teams?

Physical and business risk

IIoT security planning starts with a cyberrisk assessment. Not all IIoT deployments pose nefarious threats to the physical world. When deploying hardware that is only physically able to monitor but not control anything, we generally face only conventional business risks. Conventional enterprise security principles apply, and direct connectivity to enterprise and even cellular and Internet networks is appropriate.

For example, consider a system of thousands of solar-powered rainwater measurement devices distributed throughout a watershed as part of a water treatment flow prediction system. If the switches are compromised, or for that matter physically kicked under a rock by passing tourists, there are no grave consequences to the water system. The system is massively redundant, and device inputs are constantly correlated with external inputs, such as official meteorological reports of rainfall in an area.

But suppose the rainfall-monitoring devices can also control switches that are connected to, say, an irrigation system to activate or deactivate irrigation in an area based on the rainfall it receives. Now there are potential physical consequences of compromise. Worst-case physical consequences might include flooding, washouts, and physical damage to irrigation canals.

If monitor-only IIoT edge devices are connected to conventional control networks, we have a different problem. For example, what if the monitor-only rainfall sensors that are deployed inside the boundaries of a large water-treatment facility were connected to the facility's OT network? These connections exist because that water-treatment OT network is the easiest one for the IIoT sensors to access. In such an example, compromised monitor-only sensors give attackers an opportunity to pivot their attacks into the facility's control-critical network.

Microsegmentation

When unacceptable physical consequences of compromise are possible for IIoT deployments, we need strong protections for the edge devices. In these scenarios, a good place to start is microsegment control-critical sets of equipment or networks using unidirectional gateway technology.

Unidirectional gateways are described in section 9.2.6 of the Industrial Internet Consortium Industrial Internet Security Framework (https://www.iiconsortium.org/IISF.htm). These gateways are the strongest of the network segmentation options described in the framework. Unidirectional gateways provide additional protections to edge devices when endpoint protections in those devices are not sufficient. They enable safe flows of monitoring information to enterprise and cloud systems for big data analysis and other benefits, while physically preventing any information flow back into the edge devices.

Where to deploy the gateways is the question—in complex OT networks, unidirectional gateways may be deployed close to the edge devices, close to the connection to enterprise or Internet networks, or anywhere in between. What has emerged as a best practice is perhaps obvious in hindsight—enterprise security teams need to sit down with engineering teams and work out a strategy. Both teams need to agree on where to deploy at least one layer of unidirectional protections.