Understand OT Security

Lucien Sikkens (CISSP/GICSP)
Director OT Security at TCS
5/20/2024

Contents

1. Introduction

This document provides insights regarding the challenges of securing operational technology (OT) environments against cyber-attacks. It can be considered as a starting point for people and organizations that need to improve the OT security posture and want to avoid common pitfalls.

Most of all, the objective is to provide an understanding of what OT is and why it's different from IT. A proper understanding of OT and its peculiarities is the minimum requirement for running any successful security improvement project in OT environments.

The document is based on the author's experience in securing OT environments in various sectors. It starts with explaining OT and addressing the differences between IT and OT, followed by recommendations for successful OT security improvement programs and some real-life anecdotes illustrating practical cases.

Although quite a few subjects are addressed, please bear in mind that this document does not detail every potential subject or security control.

2. What is Operational Technology (OT)?

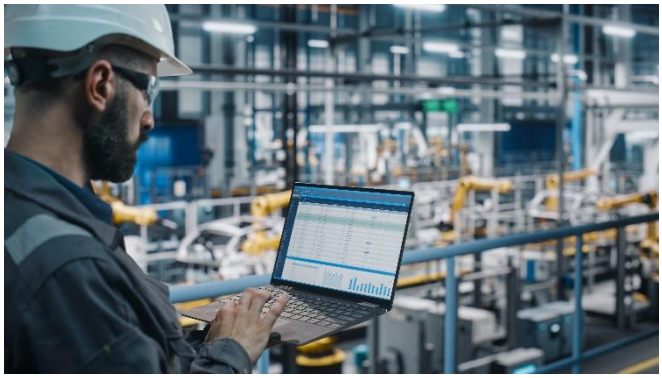

Where Information Technology (IT) is all about the automation of information, Operational Technology (OT) is all about the automation of physical processes. A couple of years ago, it was strongly tied to the industrial sector and its modernization with Industry 4.0. Before OT, we therefore often referred to it as Industrial Control Systems (ICS) or Industrial Automation & Controls Systems (IACS). At the same time, in the oil and gas sector, it can be addressed as the Process Control Domain (PCD). Nowadays, we generally use OT to cover all automation of physical processes, as we find OT in a much broader field than only industry and oil & gas. OT is all around us and applied in more than 60% of all sectors.

Although our lives have come to depend on OT, most people still don't know what the abbreviation stands for. Unlike IT, where we have been made very aware of its influence on our lives. Information Technology has automated information processing, and we are quite conscious of our interactions supported by IT. We use our computers, smartphones, and tablets to communicate and transact with banks and government agencies, order products, and book our travel.
We are much less conscious about OT, which provides us with energy and drinking water, transports us by elevators, cars, ships, and planes, and plays a vital role in the production of food, medicines, and other products we use in our daily lives. OT even keeps us alive if we think about the ventilators used in hospitals and a variety of medical machines. The impact of a successful cyber-attack on OT can be much more significant than attacks on an IT environment and requires our awareness.

3. The current attention for securing OT

The resilience of OT against cyber-attacks is generally very low, while the impact of a successful cyber-attack on OT is, in most cases, very high. This is nothing new and has been the case since the automation of OT, but vulnerabilities were extremely difficult to exploit because the access to OT was

well protected. Physical access to OT was required, and facilities were adequately protected by physical access control and social control, as small groups of operators all knew each other.

Further automation and business needs have resulted in the interconnection of IT and OT, which has opened logical doors to OT environments that haven't been prepared for such access.

The pace of business-driven convergence between IT and OT and the available budgets to reach has been much higher than initiatives and budgets to secure OT in advance properly, or at least in parallel with these initiatives. This had multiple reasons:

- Ignorance of the security posture of OT, its vulnerabilities, and resulting risks
- Resources availability for IT/OT convergence vs for securing OT
- Internal governance between business, IT, OT and security
- Lack of knowledge and understanding of OT from a security perspective

In 2010, an Iranian nuclear facility was hit by a well-known cyber-attack. The Stuxnet worm crippled the production facility and even physically destroyed nuclear enrichment centrifuges by small changes to their speed, causing them to become unstable. This attack shook the OT world, proving cyber malware could have a disastrous effect on OT.

Although various organizations used this wake-up call to define standards and recommendations for properly securing OT against cyber-attacks, most OT-driven organizations still didn't act. OT cyber-attacks were still rare news items in those days and why should anyone target them?

The past years have shown that cybercriminals do have an interest in OT environments and have built the capability to attack them. OT is particularly interesting for malicious state actors, trying to attack the critical infrastructures of nations. The Russia-Ukraine war has every National Cyber Security Centre (NCSC) on high alert, and the number of attacks on OT has drastically increased ever since.

3.1 Compliance

A clear driver for organizations to take action is compliance and regulations. For Europe, the NIS2 should become active in the last quarter of 2024, although local implementations by member states might take some more time.

On a high level, one might say that the NIS2 follows the same implementation process as the GDPR in 2016. Where GDPR focuses on data protection and privacy with the assurance of confidentiality and applying the need to know/need to have principles, the NIS2 will cover breaches to logical IT and OT networks in general on all security aspects, like safety, availability, integrity, and confidentiality (see 5.1).

The NIS2 forces organizations to take responsibility by applying fines when negligent. It not only stimulates organizations to improve their own security posture but also to pay proper attention to the area of vendor management from a security perspective, reducing the risk of supply chain attacks.

3.2 Cyber Insurance

Another stimulation comes from the opportunity to reduce cyber insurance premiums. Insurance companies have learned over the years that the average security posture of OT environments is low and have increased their premiums. Often, premiums can be reduced

when an organization can prove to have taken adequate security measures, reducing the likelihood and impact of successful cyber-attacks.

Some organizations have experienced cyber insurance not covering damages in certain conditions, for example, when the cyberattack turned out to be an act of war. This was the case with the NotPetya malware in 2017, as its origin led back to a state actor. Most organizations that were hit by NotPetya were not targeted but collateral damage by the supply chain.

4. OT security expertise is scarce

The demand for OT security specialists who can help reach an adequate OT security posture without disrupting the OT environment by their own efforts has grown significantly over the last two years, while the number of available specialists has stayed low.

One way of measuring this is by comparing the number of specialists that are Certified Information Systems Security Professionals (CISSP) versus the number of certified Global Industrial Cyber Security Professionals (GICSP). In 2022, there were 156.000 CISSPs globally versus only 4.000 GICSPs. A similar comparison could be made between security specialists who are knowledgeable about ISO 27001 versus IEC62443, which is the current standard for OT and has come forth from the ISA99.

As most universities don't offer specific OT security education, the gap continues to grow.

IT security specialists can be trained for OT under the condition that they are willing to adapt to the culture and peculiarities of the OT environment and accept that standard solutions used in IT might not work in OT or even are disruptive on their own.

Another good way to increase the capacity of OT security specialists is to train OT engineers in the field of cyber security. Engineers bring valuable expertise to the table regarding the systems used in OT.

Due to the need for OT security specialists, organizations should strive for mixed teams with regard to seniority and field experience.

5. The characteristics of an OT environment

Anyone with the ambition to improve cyber security in OT first needs to understand the characteristics and culture of OT and how they differ from IT. It's like learning another country's language, culture, and habits in order to build successful business relationships.

An important step in initiatives to improve the security of OT is gaining trust from the organization's stakeholders on the OT side. It's good to realize these stakeholders might have lost their trust in IT and cyber-security specialists by negative experiences in the past. Attempts to secure and improve the OT environment might have been made with the best intentions but could have resulted in disruptions and costly downtime. The objective of OT is to avoid any risk that might have a negative effect on safety and availability or production capacity. OT stakeholders will, therefore, be very cautious towards any change in their environment and need to build up the trust that initiatives to improve the security of their environment will take this into serious consideration, with a proper understanding of OT and how to minimize the risk of production loss.

5.1 Cyber security and the CIA triad

In cyber security for IT, we use the CIA triad to protect the confidentiality, integrity, and availability of information and information systems. Although debatable, the priority is often in this same order. Confidentiality has always been key in IT, enforced by regulations like GDPR, and given a lot of attention in governmental campaigns and corporate security awareness training. Maintaining data integrity is a strong second, as we need to ensure our ability to rely on the information provided. The importance of availability has grown over the years. Outages have been minimized with the 24/7 economy but are still more considered a nuisance when they happen.

In OT, we use the same aspects for security but with reversed priority. More importantly, human and environmental safety is added as the highest priority. We expect OT to function continuously with the highest efficiency and minimal downtime. Think of the last time tap water wasn't available. Most probably quite some time ago and experienced as more disruptive than downtime of IT.

Like in IT, integrity is a strong second, assuring product quality and systems reliability. For that matter, "A" and "I" go hand in hand. Coming back to Stuxnet, affecting integrity caused a loss of availability.

Confidentiality is not so much an aspect of OT. One might think that the recipe for a beverage or medicine should be kept confidential, but those recipes are stored in the IT environment, while OT uses just its data without context. Only with complex reverse engineering might those recipes be unraveled.

In IT, we use CIA, and in OT, we use SAI(C), which addresses Safety, Availability, Integrity, and Confidentiality.

An adversary only needs one flaw

While discussing an OT security assessment with a client, I explained the concept of red teaming. This particular client had implemented rigorous physical access control as it was required to take anti-terrorism measures. "You'll never get in!"

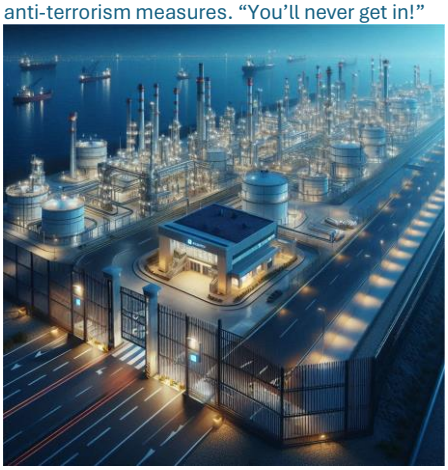

Red teaming is one of the most challenging and interesting tasks a security team can do and it stimulates creativity as we step into the shoes of an adversary. We agreed to a period of 4 weeks in which we would try and breach security, with the objective to physically enter the control room. Reconnaissance started and we felt the client might be right in their bold statement. The guards at the gate performed their work very well and the place was covered with CCTV, high fences and only one entry. Then we discovered something interesting. Most employees needed to swipe their badge at the gate, but for few the gate opened by recognizing the number plate of the car. A blank numberplate was acquired, self-stickered with the numbers of a plate we photographed and mounted to the front of our car. Adrenaline rushed when driving up to the gate and..... it opened! On the parking lot a friendly employee escorted our team, without questions asked, into the control room, as they said they needed to do maintenance there. All employees were very much aware of the rigorous physical security measures and assumed that anyone on the premises was there for a valid reason. It turned out that the Automated Number Plate recognition had a flaw as well, as the car with the valid number plate passed the gate earlier that day.

5.2 OT is (in)secure by design

The Stuxnet malware that crippled the Iranian nuclear enrichment plant opened the eyes of the world regarding the vulnerability of OT environments. This example is still used in discussions about improving the cyber resilience of OT environments and invigorated by statements like "OT is insecure by design."

That statement, however, is completely incorrect, as we can deduct from the previous paragraph.

OT environments have been designed to be extremely high in availability and reliability. This has been ensured by drastically reducing the latency in OT network communication. This latency reduction has been achieved by, for example, removing overheads like encryption and authentication from network communication. This does make logical access cyber-insecure, but that was remediated

by only allowing access within the physical environment and implementing strong physical access controls. Although Stuxnet proved that strong physical access control could fail, from a risk management perspective, the likelihood is very low as it would require the adversary to be locally present and able to breach both the physical access controls and social control.

When most OT environments were designed, they were designed with security in mind for the intended use at that time. It's like a ten-story building designed to provide safe housing to its inhabitants. Build another 20 stories on top of it; the construction might not hold and become unsafe. In OT, we relied on the protection of physical access gates and doors but have built logical doors that provide access from the IT- and even the Internet. Doors that can't be seen and aren't monitored very well.

5.3 Cultural behavior

The biggest pitfall for programs that aim to improve the cyber resilience of OT environments is failing governance and misalignment of stakeholders, resulting from the absence of mutual understanding. It's a people's challenge in an OT world driven by people, processes, and technology.

Some things are easy to spot, like the white collars in the office environment and the blue collars in OT. This is immediately tied to the priority of security objectives. Starting with safety, clothes worn in OT are important in assuring just that. People in OT will use the stair railing, while we don't see the same discipline from office employees. Again, ensuring personal safety, like parking a car backward so one has a better view and can leave quicker in an emergency.

People think "logically" in IT, while in OT, people rely more on their physical senses. They act upon what they see, feel, hear, smell, or taste and often must react swiftly when safety or production can be affected.

An interesting example from a security assessment in a big factory showed how people interpret signals. The plant manager confidently pointed to a device in his server room and said his factory was secure. When asked for an explanation, the answer was, "The light is green." The device pointed at was the logical firewall between the IT and OT environment, and the light indicated that the device was powered up and running. After examination, it turned out that too many logical doors in that firewall were left open, making the device useless from a cybersecurity perspective. The state of those logical doors was something the plant manager couldn't see, unlike the state of the physical doors in his facility and the green lights at every machine.
In OT, a green light means safety is assured, an amber light indicates an issue and a red light indicates danger.
A firewall vendor recently adopted this common signaling and way of working by implementing I/O (input/output) functionality. With the turn of a physical key, a logical door (VPN connection) could be opened, and the state of the door was reflected by a RAG (red-amber-green) tower light, which indicated in red that the door was open and in use.

Another example comes from a pharmaceutical company. They had strict procedures for allowing suppliers to perform maintenance of their systems in the OT environment. The engineer needed to report at the reception, providing an ID and was then physically guided to the systems that needed maintenance. During the maintenance, the engineer was observed, and it was made sure that the engineer was escorted out when done. For efficiency, the company also allowed remote access and gave "the key to the environment" without knowing who was using it, when, why, how long, and what was done. Access via a logical door that couldn't be seen through their eyes.

5.4 Language

Language goes hand in hand with culture; we need to speak each other's or a mutual language to understand. As in any discipline, skilled and experienced people start using abbreviations, which doesn't make life easier for those unfamiliar with the environment.
A good example is the HMI. People in IT work daily with them, but no one will use this abbreviation, while it's one of the most common in OT. The Human-Machine Interface, which can take the form of a touch screen with a GUI (graphical user interface) in OT and run on Windows XP as OS (operating system), is very similar to working with computers and tablets in IT.
A PLC is a Programmable Logic Controller, or a very small computer, that acts upon input from sensors and controls actuators like a motor. You would find them in many places without knowing they are there, and PLCs are, for example, responsible for stopping an elevator on the chosen floor. Operators of OT need to monitor and control production processes and, therefore, use Supervisory Control And Data Acquisition systems, or SCADA. And yes, an operator can interact with SCADA using the HMI. SCADA systems are used for single sites or processes, while Distributed Control Systems do the same for complex processes and multi-site environments.

Knowing and understanding the abbreviations and technology make communicating with the stakeholders of an OT environment and gaining their trust easier.

5.5 Policies and procedures

OT relies on the adequate implementation of safety measures to keep people and the environment safe. Policies and procedures are readily available from a safety perspective but might be lacking from the cyber security perspective or are copies of IT policies.

Being able to show the availability of policies and procedures might be considered a tick in the box, but more important is whether they are implemented and used, as the documents themselves won't keep the environment safe.

Implementation means creating the ability to apply the policies and procedures adequately. Due to the cultural differences between IT and OT, there is a need to tune the content to the intended audience and characteristics of the environment. On a higher abstraction level, the essence of policies and procedures might be the same for IT and OT, but they can differ operationally when implemented. It's comparable to the creation of a Word document versus a PowerPoint. The message might be the same, but the content presentation is tuned to its use.

In OT, we must act swiftly and pragmatically when something is off. Quick and unambiguous actions to assure safety and availability are guided by easy-to-read and understand policies and procedures, often one-pagers with symbols instead of lengthy Word documents.

An example of where policies and procedures differ in OT from IT is patching. On a high abstraction level, IT and OT will share the same objectives of keeping the environment safe and secure to assure safety, availability, integrity, and confidentiality. Patching vulnerabilities is, for both environments, a best practice to be applied. However, the policies and procedures for patching can differ between the two environments regarding the frequency and urgency of patching. For an IT environment, we might have decided to implement a critical patch within one week after it became available. Still, for OT, we might decide to implement the patch much later. In both cases, this is all related to risk management and risk appetite as determined by the business. Patching in OT will often cause downtime, resulting in production and financial loss. This loss should be weighed against the potential loss incurred by a cyber attack exploiting the

vulnerability that needs to be patched due to the likelihood of a successful attack and its impact. If other security controls are in place that reduce either or both likelihood and impact, the calculated loss from a successful cyber-attack would be decreased to an adequate level. In that case, the patch can be implemented at a later stage.

5.6 Cyber security awareness in OT

As a result of the cultural aspects of OT and focus on events that can be observed with one's senses, cyber security awareness is often at a lower level than in the office IT environment. Cyber security awareness training has generally improved over the years but is still very much tuned to the processes and cyber threats that affect IT. Modern awareness training can use gamification, which is brought to the desktop or laptop of office workers. At the same time, in OT, the operators might not even have such systems or personal accounts to which the content can be delivered.

The difference between safety and cyber security awareness in OT is visible in many organizations if we look at procedure charts for calling in a safety issue. These charts are spread around the OT environments with phone numbers that are easy to remember in case of emergency. Rarely similar charts are available for reporting cyber incidents, and in most cases, the emergency desk will be challenged in handling reported cyber incidents.

It's quite common in OT for employees to watch a safety video and answer related questions about the content before being physically allowed on the premises. This does not only cover the organization's employees but also contractors and vendors making deliveries or providing maintenance.
These videos address threats to personal and environmental safety but, in most cases, don't include cyber scenarios. The delivery method of the awareness training content, using video material created on-site, can also be used for cyber awareness training customized to the organization instead of commercial off-the-shelf awareness training.

5.7 Depreciation time

An important aspect of OT to consider is the long asset depreciation time in the OT environment.

In IT, the depreciation time is short. We replace assets like smartphones, tablets, laptops, and servers between once every two years and once every five years.

Due to the sheer number of IT assets, hardware and software vendors have adapted to the depreciation cycle and limited their support. The support period for modern operating systems is a good example. Over time, newly discovered vulnerabilities, potentially caused by new techniques and capabilities of adversaries, will no longer be fixed as outdated operating systems are expected to be replaced.

OT assets have been built to run permanently and can last 25 to 40 years or, incidentally, even longer. The formal lifespan of a nuclear power plant is, for example, 30 years, similar to that of a Boeing 747.

OT systems launched in 2002, running Windows XP, were considered state-of-the-art then but might still be around in 2032. This would be 18 years after Microsoft ended its already extended support.

The main reason for the long depreciation time is cost. The financial impact of replacing a laptop is minimal, and a new one can increase efficiency. Efficiency and security justify the investment.

The cost of assets in OT is incomparably higher and is often part of a bigger system built to last. Unless the whole production process is redesigned and all assets are replaced, a single asset replacement will often not gain increased production. In addition, it might cause a compliance issue for systems that have been validated.
Thus, replacement comes at a high cost, plus the additional cost of downtime, loss of production, and potential re-validation. In most cases, the business case for early replacements is negative due to security requirements. It must be noted that cyber vulnerabilities might be underestimated in the equation under the assumption that the OT environment is still adequately protected by physical security measures, resulting in a small attack surface.

Another aspect of depreciation is dependency. Although a simple workstation running Windows XP in a factory might seem easy to replace with a new workstation running Windows 11, the software required to program PLCs might not be available for the modern operating system. Replacement of such a workstation might be IT's best practice and remediation of the vulnerabilities Windows XP has. Still, it causes a self-inflicted denial of service for production regarding PLC programming.

5.8 From legacy to IIoT

Newly built OT environments can be state-of-the-art, but most of them are legacy due to the aforementioned depreciation time of assets. Even in newly built environments, we already find assets that we consider legacy, viewed from an IT perspective. New sites can be designed to use "proven technology," meaning that the technology has already existed for at least a couple of years.

On the other hand, organizations need to stay competitive and meet regulations, which require changes to their OT environments. Gaining data from the production environment is the first step in evaluating and tuning production performance, improving customer information about production and delivery times, and providing evidence for regulators and auditors.

The first logical step in making data available from OT to IT is providing access to the historian, that maintains production data records. This means that either access to the historian server needs to be provided, or its data needs to be replicated. This should lead to a decision regarding the network architecture and design principles that will maintain an adequate security posture, as a newly implemented "logical door" could have undesired consequences.

More data means more control, and data can be provided by sensors. Sensors nowadays have become smaller smarter, and can be easily added to existing production systems. Ideally, they use wireless communication and potentially the Internet to share their data wherever required. These sensors have become part of the Internet of Things, and while we talk about OT and the production industry, let's call it the Industrial Internet of Things (IIoT).

Interestingly enough, modern IoT has learned lessons from OT regarding efficient network communication. It uses Message Queuing Telemetry Transport (MQTT) as a lightweight network protocol that reduces the network overhead. However, although we regard IoT as modern, MQTT was already developed 25 years ago.

In essence, we have created an OT environment that is a mix of old and brand-new equipment, with various technologies getting increasingly entangled. Maintaining complete insights regarding all assets, their roles in the network, and how they can influence production is a challenge. IIoT has moved from nice to have to critical, and threats, like data manipulation, should be considered data that could be used to control elements in production.

5.9 OT Asset visibility

Most organizations that rely on OT do not have complete, documented insight into their OT assets. In practice, visibility varies between 40% and 80% due to poor asset management. This creates a big challenge for projects that aim to improve the security posture because parts of the foundation are missing. You can't protect what you don't know.

A best practice for securing an OT environment is the implementation of network segregation and segmentation. Still, to group assets in a segment, one should at least know the assets that need to be grouped.

There are multiple reasons for not having 100% asset visibility. Due to the long depreciation time and the need to modernize, OT environments grew over time. The oldest asset might have been there for over 30 years, while the newest asset has just been implemented. With the pragmatic culture and focus on high availability, assets have been added to the environment with short implementation time and no proper procedures for registering the assets. Knowledge of the systems is often available at the shop floor level and in people's minds, but organizations are challenged with this knowledge diluting when people leave or retire.

Manual effort can gain asset visibility, but it can be time-consuming. A pharmaceutical company made an educated guess and estimated one labor hour per asset to be discovered and documented in an environment where they assumed about 2000 assets to be present. This would result in a 1,25 FTE year. Fulfilling the prerequisite of a network segmentation project takes a long time.

Shodan.io

Lacking asset visibility is a problem in OT environments, but it will even be a bigger problem when adversaries have insights you're not aware off and potentially regarding specific assets that are missing in an organizations view.

In client conversations I might use the capabilities of Shodan and will almost always find that the tool is completely unknown, while everyone knows Google. Where Google answers our questions for information (IT), Shodan informs us about systems connected to the Internet. The Internet of Things! You might find a variety of badly secured cameras at peoples homes, but also remote connections to wind farms, water treatment plants and things you could not even imagine were Internet connected.

Enhance your queries on Shodan and correlate data from the Automatic Identification System (AIS) used on ships with common communication and control equipment on board and you might find a bulk carrier that you can gain access to. Shodan your own organization and you might discover some of your systems. Not to worry immediately, because Shodan identifies IoT and your systems are probably intentional connected to the Internet. But do check whether access has been secured adequately.

In IT, scanning solutions are used to discover assets electronically. This is a no-go for OT as such scanning solutions might be disruptive for certain OT assets. Even a simple ping command could cause the asset to behave undesirably, potentially affecting production. However, a scanning solution would certainly contribute to the efficiency of asset identification and can provide additional information about vulnerabilities. It requires specific tooling that can guarantee non-intrusiveness in the OT environment, and one needs to realize that manual effort will always be required to record metadata that is not electronically discoverable by such tools. The previously mentioned pharmaceutical company reduced the duration for gaining full asset visibility to only six weeks in this way.

The next question is how to store the asset information. Although many organizations use modern configuration management databases (CMDB) for IT, many also use MS Excel, MS Access, or similar for OT assets. Before deciding to use IT's CMDB for OT assets, one must first get clarity on the

requirements for administrating OT assets, as they can differ from IT. OT might require specific metadata to be administered, and security and trust can play an important role. Sensitive data regarding OT assets might be stored, and access to this data is limited to people on a need-to-have basis.

5.10 Architecture and the Purdue model

The Purdue model is a best practice for OT network architecture and design. It is well explained in the IEC62443 standard and used by organizations to create a reference architecture and implement segregation and segmentation to reduce the cyber attack surface and propagation of an attack on the network.

Let's start by explaining what the Purdue model is. It's a layered architecture model in which we distinguish the technologies used in OT and apply appropriate security controls to the different layers.

Cloud and Internet form level 5, and enterprise IT is covered in level 4, which can be considered the tip of the iceberg.

Below the surface of the waterline, we find the OT environment, with central site operations and control on level 3, including the Manufacturing Execution Systems (MES) in a factory. In level 2, we find local supervisory control, where operators control the process using their HMIs. Level 1 contains the basic input/output (I/O) controls covered in programmable logic controllers that act upon sensors in level 0 and control the actuators, like motors, based on the sensor value.

We create level 3.5 on the surface of the waterline, also known as the demilitarized zone (DMZ) between IT and OT. This zone typically holds central management systems and systems that provide data for business purposes coming from the OT environment, like historians.

6. Where to start

6.1 Awareness

Realizing that one's organization is highly dependent on OT for achieving business results and that OT has become more vulnerable over time is the first step in improving the cyber resilience of the OT environment. This awareness can come from having already suffered a successful cyber-attack, from cyber-attacks on similar organizations published in the media, from new insights, consultancy, and/or regulations.

There's no easy fix or off-the-shelf solution to secure an OT environment adequately. It will take an investment to set up an improvement program that might run for many years at significant costs. The art here is to spend the money wisely, acting smart on priorities and understanding the value of required and desired security controls.

Tabletop

A tabletop exercise is a good way to increase awareness and test the effectiveness of communication and actions while applying the business continuity and disaster recovery plan.
I have always put quite some effort in the creation of tabletops as they should be as realistic as possible. Basically, it's similar as the reconnaissance phase of a red teaming engagement by performing research on an organization's vulnerabilities and creating attack scenarios.
For an organization in the maritime sector, the idea was to disrupt or take over the control of their vessels. Research showed that a supply chain attack could be used to achieve this objective and these findings were validated on paper. A scenario was built and unfolded during the tabletop exercise, raising eyebrows of the participants. The scenario was initially perceived as farfetched, but when the underlying proof was presented the awareness that such an attack could happen was immediately there.

The overall budget for cyber security is usually around 2~3% of an organization's revenue, but remember that most OT-dependent organizations are significantly lagging on their OT security

posture. The bottom line is that a serious budget is required, and it needs to be managed well to avoid overspending.

6.2 Governance

Understanding that the organization is facing a long-term security improvement program that requires a significant budget calls for a good governance structure to prevent failure.

It requires the knowledge and experience of all relevant organizational stakeholders, supported by the board level. Relevant stakeholders will come from business, operations, IT security, and potentially legal disciplines.

Initially, the group of stakeholders can be small, covering tasks like performing a business impact analysis, acquiring more insights, defining the program, objectives, strategy, and roadmap, selecting standards, identifying compliance requirements, and developing principles like reference architecture and design principles.

After the foundation has been laid, the program can kick off with defined projects, and the number of stakeholders will increase, focusing on their individual projects. The project scope will be related to the strategy and priorities and can relate to sites in a multi-site organization, production lines, or technology/security control to be implemented.

6.3 The Business Impact Analysis

From a risk management perspective, the cost of improving and maintaining the security stature of the OT environment should be less than the expected loss resulting from one or more cyber-attacks on the environment in a period.

The calculation of loss is hypothetical and goes hand in hand with the organization's risk appetite, although regulations might force the implementation of certain security controls. For that matter, it’s like having insurance with deductibles.

An organization with multiple production plants should be able to accurately calculate the financial loss of downtime on a per-plant basis. Suppose the downtime is a result of a cyber-attack. In that case, the duration can be predicted from variables like tested recovery point (RPO), recovery time objectives (RTO), and market experience with similar attacks. As cyber-attacks are not tested in production environments, the recovery time can significantly deviate from the RTO. The total loss at a cyber attack will not be limited to missing revenue due to downtime but will include recovery costs and penalties from regulators and the effects of reputational loss.

The likelihood of a successful cyber-attack occurring is even more difficult to determine. However, we can put things into perspective from an OT point of view. One thing that has drastically changed is the number of attempts to get in. In the old days, plants were protected with gates and fences, and a guard monitored all access, preferably on one central access road. A criminal adversary had to be physically at the location for an attempt to break in. Nowadays, with so many logical doors connected to the Internet, adversaries can come from all over the world with various backgrounds. From any kid with a computer rattling at a logical door to state actors with unlimited resources. It’s a fact that those logical doors are rattled continuously. The question is what they provide access to and what harm can be done. The likelihood of suffering great loss due to a cyber-attack is much higher when preventative security controls are lacking.

A business impact analysis can be conducted on various levels. From the enterprise level to the plant, production line, and asset level. The impact can be determined from individual aspects, summing up to an overall impact rating. Together with the likelihood, the result is the well-known format of a risk matrix, as shown below.

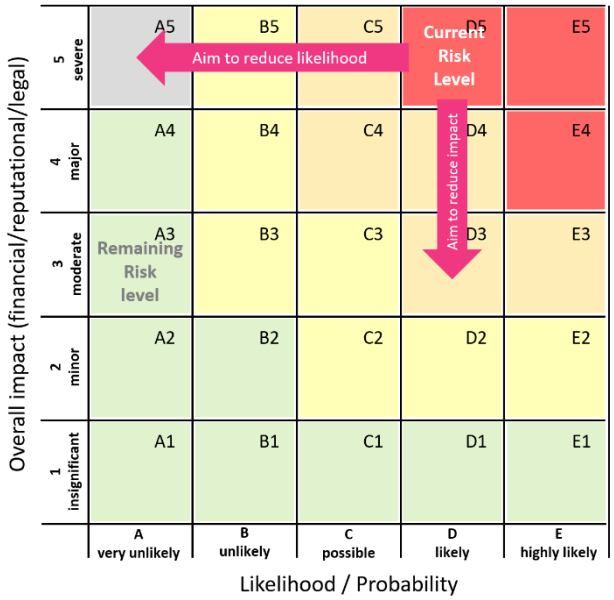

In this diagram, the impact of a risk occurring is estimated as severe and the probability of the risk occurring is likely, resulting in a risk score of D5 in the matrix.

The objective is to lower the risk score by taking measures that reduce the impact when the risk occurs and reduces the likelihood of the risk even occurring.

These are the two vectors than can also be related to specific security controls.

The vectors reduce the risk from the D5 score (red) to an A3 score (green) in the given example. An organization can determine the target score and implement security controls that result in this score or implement security controls and evaluate whether the resulting score is acceptable.

On a more detailed level, the risk can be related to the perspective of what we want to secure, meaning safety, availability, integrity and confidentiality. When performing a business impact analysis on that level in an OT environment, the impact on confidentiality is often low for most systems when they don't store or process confidential information. From a production perspective a system might have a high BIA score on availability. With this detailed level, a score might look like 3-5-3-1 on SAI(C), where 1 is low and 5 is high.

6.4 Vectorizing security controls

The objective is to adequately reduce the likelihood and impact of a risk occurring and implement the security controls that will contribute to that.

Security controls can be mapped in the risk matrix, using vectors to influence impact and likelihood reduction. Vectorization can help determine the priority for implementing specific security controls.

1 Security control 1 could be encryption of data. When the risk is defined as confidential data becoming public as the result of a cyber-attack, encryption will reduce the impact but not the likelihood of data being stolen. It therefore has just one vector.

2 Security control 2 could be the implementation of rigorous access control. This will reduce the likelihood an attacker being able to access the data, but will not reduce the impact when the access control is breached.

3 Security control 3 could be segmentation, reducing both likelihood and impact of a cyber breach. As explained in 6.7, segmentation will divide a network in multiple smaller networks. If one smaller network is breached, the impact is smaller than the impact for the sum of all networks. The level of segmentation will determine the length of the

vector, reducing the impact. Segmentation will also increase the network complexity for a hacker, so the likelihood of all segments being breached is smaller than when there would be no segmentation at all.

The example on the left might show the risk level regarding the production of an organization being affected by a malware attack on the OT environment, crippling the production lines. Not remediating the vulnerabilities in the OT network that drive this risk to a D5 level will lead to the risk even increasing over time as capabilities of adversaries grow.

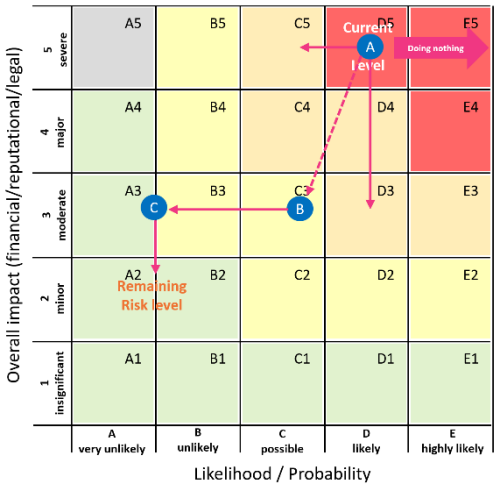

Implementing segmentation (A) will reduce impact and likelihood. In this hypothetical case to a level C3. Rigorous access control (B) will further decrease the likelihood to A3/B3 and improving backup and restore will decrease recovery time and thus impact to A2/B2.

The order of security controls to be implemented, the level of implementation (for example level of segmentation) and pace of implementation will determine the path coming from D5 to A2/B2. Changing the order of the vectors or doing things in parallel will change the path and having insights in these vectors can help in making decisions with regards to priorities.

6.5 Deeper insights

Before starting a cybersecurity improvement program for the OT environment, it is wise to assess and determine the situation. Assessments cover organizational maturity regarding cyber security, the state of security versus the IEC62443 standard, and a network analysis covering people, processes, and technology. It's a worthwhile investment as it identifies the main vulnerabilities and helps prioritize and plan.

In most cases, the maturity and IEC62443 assessments are executed by qualified consultants using a questionnaire and a scorecard. The consultant can provide context to the questions and will ideally combine the interviews of selected stakeholders with a validation and discovery phase. In this phase, the consultant will physically examine the OT environment, checking and validating the answers that were given during the interviews and observing and discovering relevant but unresolved issues.

An example lies in policies and procedures. An organization might reply in an interview to have these documents available. It can even show their existence, but during validation on-site, the finding can be that the policies and procedures are not known by the intended users, making this security control ineffective. Another example is access control. Physical access control and user accounts with a proper password policy might be in place. Still, if doors are held open by wedges and passwords are noted on sticky papers on monitors, the intended security controls will become ineffective.

Instead of interviews, self-assessment questionnaires can be used. Without a consultant providing the context of the questions, this will only work well when that context is provided by the system holding the questionnaire. The advantage of questionnaires is that they are more efficient in planning answers and lowering costs. Ideally, a specialist can be consulted for

specific questions with doubts or contradictory answers from multiple stakeholders of the organization.

Network analysis provides great insights about assets in the OT environment and is highly recommended to be executed. An absolute requirement of such an analysis is that it should be conducted non-intrusively. This requires a specialized solution that scans passively while understanding the various protocols in an OT environment and properly identifying the assets. Some solutions can provide additional information after accurately determining the assets and understanding how additional information can be polled without being disruptive. Often, organizations will not go that far due to a lack of trust, and some will even decide to use data diodes to guarantee the non-intrusiveness of the scan.

A clear value of the network analysis is increasing the visibility of assets, as many organizations lack 20 to 60 percent visibility. A network analysis can quickly bring that visibility up to over 90%. However, one needs to realize that the analysis can only identify assets communicating on the network during the scanning period. In addition, the analysis will only show information that can be deducted from this communication, and certain desired metadata, like, for example, the owner of the asset, will not be part of that. Therefore, network analysis is often accompanied by manual labor to complete the last visibility percentages and add additional metadata. Still, it has been proven that a combination of network analysis and added manual labor will be far more efficient than doing it all by hand.

In addition to asset visibility, the analysis will also provide information about the communication between assets, protocols used, external connections, dual-homed systems, asset vulnerabilities, and so on. This is valuable information for network segmentation and hardening activities.

6.6 Improvement program

With the previous four steps covered, a cyber security improvement program can be set up. The business impact analysis will be used together with the insights obtained to determine priorities, the required budget, and the estimated duration.

An improvement program can span multiple years, and the objective for the 1st phase is to realize the biggest risk reduction in parallel executing the easy fixes.
Most organizations will start with:

- Implement or improve policies and procedures and execute awareness training
- Implementing or improving network segregation and segmentation
- Implementing or improving remote and local access control
- Implementing or improving actionable monitoring

The network analysis, as mentioned under 6.4, is of great value for the projects mentioned above. Evaluating the existing firewall capabilities and rules is worthwhile before acquiring new firewalls. Firewall vendors can offer a proof of concept, placing their well-configured firewall behind the existing one to demonstrate its ineffectiveness. However, in many cases, a firewall clean-up project can be sufficient and executed at a lower cost than acquiring new firewalls. This depends on the functionality and desired capability, as the latest firewalls can go further in monitoring and analyzing traffic and take automated action.

Under the program, multiple projects can be run in parallel and with a different focus. In general, the traditional waterfall approach will be the best fit for the OT part of an organization

unless they have matured in agile working. Realizing that the security posture of OT is significantly behind on IT and with threats growing fast, some organizations have started with an agile way of working as the holy grail but failed due to lack of experience, resulting in many initiatives started but none leading to effective improvements.

6.7 Network segregation and segmentation

Network segmentation as security control can be easily explained to people working in OT because, in essence, it's the same as physical segmentation. Physical segmentation can, for example, be implemented to prevent a fire from spreading or, in a biological lab environment, a virus from spreading. In case of a fire, physical segmentation will decrease the impact. That's the same for logical segmentation, where the network is split up into multiple smaller networks with limited communication, similar to doors that can be shut in physical rooms.

When executing security assessments in OT, one can encounter various network setups. For organizations having a flat network for IT and OT where everything is connected, IT/OT networks that were intended to be segregated by the implementation of a firewall, but where over time logical doors have been opened for convenience, testing, or business needs but never have been closed again. Implementation of any-any rules in firewalls that reduced their functionality to an expensive switch, up to networks that have been segmented well. In most cases, though, there's quite some work on proper segmentation.

The first step is to create reference architecture and design principles. These form the cookbook for the actual implementation, assuring a standard approach.
Use the IEC62443 standard as a baseline and tune the standard Purdue model to the organization's specifics. The technical design principles are already present in IEC62443 and can be used for the organization.

What needs to be added are the business-driven design principles related to the organization's risk management and risk appetite. Let's take ten identical production lines as an example. It's good practice to segment these production lines from an availability perspective in case of a successful cyber-attack. If the ten are separated into two networks containing five lines, 50% of production can be lost when a segment has suffered an attack. By creating five segments, each containing two production lines, only 20% of production is affected. However, complexity increases, and the same goes for network maintenance. The level of segmentation is a business decision and can be supported by business impact analysis on a lower level. In the design principles, an organization might document that an asset or production line with a business impact analysis score of 4 or 5 will be placed in a dedicated segment.

Architects will now have to create local designs based on the reference architecture, design principles, and information about the assets. Adequate asset information is required to minimize the risk of unplanned disruptions and downtime. Network analysis helps speed up the acquisition and communication of asset information in the network, but it cannot be relied upon alone. Network analysis provides only a snapshot that can span some hours to weeks, but any communication outside this timeframe will not be captured and made available. Be aware that some assets might require them to validate their license periodically to avoid stopping working or accessing a remote library for data validation. Such communication requirements need to be provided by the asset owner or operator.

Once all data is available, a change form can be completed, including all required changes, a checklist, and a test approach. It forms a step-by-step manual for applying all necessary changes to create and populate the segment.

6.8 Access control and monitoring

Improving cyber security in OT is no rocket science. It just requires another way of thinking, and it helps OT staff translate the logic from cyberspace to the physical world they know. Factories and plants have been kept safe by the implementation of physical access control and monitoring, starting with the guard at the gate who knew the factory workers and, over time, even created awareness of their behavior.

The logical world of cyberspace isn't much different, so it's as simple as implementing the successful security principles from the physical world in the logical one.

Access is the key element here, and a best practice is to segregate IT from OT by implementing the logical doors of a firewall, like the physical access doors between the environments.
The next step is segmentation, which can be applied in both the physical and logical worlds.
In practice, we might find different buildings on a site, each housing an individual production line, while all the production lines on site are interconnected in one network. Segregating them into different logical networks reduces the risk of cyber malware attacks spreading over the production lines as the different buildings would protect from a fire spreading.

The same principles must be applied to remote access, as shown in the examples in 5.3.
Remote access to the OT environment can come from the internal enterprise network (level 4 in the Purdue model) or the outside (level 5). In all cases, access should only be provided on a need-to-have basis with the ability to identify the user. When providing physical access, the guard can check the ID card and inspect what is brought in or out of the environment. With logical access, we strive to do the same and to translate the logical world to the physical; it's useful to understand the security controls of the logical world and their constraints.

A VPN connection is a good method to provide remote logical access to a site, but it has limitations. It can be regarded as a key, and anyone with that key has access. Even beyond that, this key can enable a tunnel that can provide access to multiple users. From that perspective, a Privileged Access Management (PAM) solution comes much closer to how a physical guard would provide access to an individual.

COWs

The great thing of working in OT with a variety of clients is that you continue to learn about production processes, technology, habits, and language.

So, while I was performing a security assessment in a hospital, the cows were mentioned. And although the hospital was in a rural area, I couldn't imagine that they were referring to the animal.

A cow in a hospital turned out to be a Computer On Wheels and one was parked in a patients room. I observed that the cow was unlocked and discussed this with the doctor. The answer was clear. In case of an emergency, they needed every second to save a patient's live and felt they couldn't afford the time to login or forget the password under a stressed situation as this could cost a patients live.

I challenged the doctor, asking whether a visitor could access patient information and would even be able to make changes, like for example to the blood type. The answer was yes.

Reducing access is a very effective way of reducing the likelihood and impact of getting affected by a successful cyber-attack. Note that reduced access goes from humans to machines and from machines to machines.

The physical guard and fellow workers can observe the deviant behavior of a colleague or spot an intrusion while watching the CCTV system. As addressed in 5.3, they use their visual senses. Deviant behavior can have many causes and will be harmless in most cases, but it could affect the safety of an environment or other aspects like availability, integrity, and confidentiality. If there is an indication of danger, action can be taken.

Monitoring in the logical world works the same way, and with the available computing power, it will be even more efficient. It's often called User and Entity Behavior Analysis (UEBA) and applies machine learning to identify anomalies. Compared to the IT environment, this type of monitoring will be even more effective in OT as the processes in OT are more standardized and predictable, with less influence from us humans. Like CCTV for physical access, logical monitoring should be used for logical access at a minimum and has even more value when implemented on a broader scale.
Stuxnet (see 3) could have been discovered with the implementation of monitoring, as it would have indicated that the centrifuge's speed was tampered with.

Monitoring can also provide value beyond improving security, such as preventative maintenance. It can measure a motor's increased power consumption while maintaining the same speed, indicating that a bearing is worn and giving more resistance.
Monitoring, however, is only effective when it understands the language. When someone speaks Chinese, a Dutch person will hear the noise but not understand what's been said. The same goes for logical monitoring in OT networks, where not every asset speaks the TCS/IP language. This means that the monitoring solution must be able to understand the languages or protocols used for communication in OT, like Modbus, Profibus, and so on.

Like the response of the security guard, watching an intrusion on the CCTV system, logical monitoring also needs to result in a response to an incident. This can be a manual action of a response team in the (cyber) Security Operations Center or an automated action by the monitoring system. Automated responses are still rarely accepted in OT due to the risk of an automated response based on the false identification of an incident, potentially affecting production availability. Only when it's guaranteed that there will be no false

Vehicles and vessels

Although we tend to quickly think of factories and utilities when discussing OT, we must realize that OT can be found in so much more places. And all of them undergo modernization up to the level of Internet connectivity.

There are good documents describing all facets of Industry 4.0, but if we look at Shipping 4.0 the essence is almost the same.

The military are relying on command vehicles that are a combination of OT and IT on a small but very critical scale and under both physical and cyber attack during operations. In the last decade they experienced the effects from remote cyber-attacks on these vehicles, affecting the control and trust in the vehicle's systems.

Independent monitoring was implemented as a proof of concept to gather data from the vehicle's networks, sensors and actuators to spot anomalies and provide actionable information to the vehicle commander.

On of the simplest use cases was related to the vehicles position, direction, and speed where correlation between information from the vehicles canbus (OT) and GPS would indicate attempts of nearby GPS spoofing and triangulation from different command vehicles could identify the adversary's spoofing location.

positives, or the automated action will not affect production will an automated response be acceptable.

Monitoring solutions come at a cost; again, the BIA process can be used to determine the scope of monitoring. Organizations might implement complete monitoring on their most business-critical sites or production lines and a more straightforward or even no implementation at small sites with low impact. While having a monitoring solution in place, its capabilities can be used for periodic asset discovery and network analysis by feeding the solution packet capture files from the unmonitored site. This is the network analysis, as addressed in 6.4.

6.9 Zero trust

The current buzz word in security is zero trust, and one might get the impression that several vendors have a commercial off the shelf solution that can be quickly implemented and brings an environment to the highest security posture. Such a panacea doesn't exist.

Zero trust is however the best practice to pursue, and the first steps will already be achieved by implementing network segmentation, access control and monitoring.

The essence of zero trust is clear. Don't just trust anyone or anything, but always verify identity and the need to know/have principle. See the anecdote about an adversary only needing one flaw, where a critical infrastructure implemented such rigorous physical security measures to a site, that everyone on site was automatically assumed legitimate. The principles of rigorous physical access controls to a "site" were already applied during the Trojan war, that took place around 1190BC. A site might be well secured, but when an adversary gets in, everything within the environment is highly vulnerable.

Zero trust comes with maturity levels, where macro segmentation, access control and monitoring forms the basis. This is the first and well achievable step for OT environments as described in this document. In the highest maturity level we have implemented micro segmentation, continuous access validation, monitoring and response on an asset level.
The challenge for OT environments I reaching the highest maturity level as this requires full asset visibility, the ability for legacy assets to handle the zero trust principles and the increased latency.

In general we can map zero trust to the Purdue architecture model and assume a high zero trust maturity for level 5, 4 and 3.5 and a lower maturity for the levels

Zero trust in power grid OT

Imagine a bustling control center in Europe, where operators monitor a sprawling power grid that supplies energy to millions. The heart of this grid's security is a sophisticated system known as Public Key Infrastructure (PKI).
This system arms each device within the grid with a digital certificate, much like a unique, unforgeable ID card that proves its identity. This digital safeguard ensures that only verified devices can communicate, forming a robust shield against unauthorized intrusions.

This approach came to life after an incident where an update introduced a seemingly minor software bug that caused erratic behavior in several grid components. Quickly, operators isolated the issue, thanks to their segmented network setup—each segment acting as a watertight compartment, preventing the spread of potential threats. This real-world hiccup underscored the importance of their Zero Trust strategy: verify rigorously, trust sparingly, and compartmentalize diligently to enhance resilience.

The European and U.S. implementations of Zero Trust in their power grids highlight a critical shift in cybersecurity mentality. No longer is it enough to build high walls; now, we must also install sensitive, discerning systems within these walls—systems that are always watching, always verifying, and always ready to respond.

The journey toward adopting Zero Trust involves deploying new technologies and transforming how organizations perceive and handle security. It requires a perpetually skeptical mindset, rigorously checking credentials, and continuously analyzing behaviors to ensure the safety and reliability of critical infrastructure.

By embracing these principles, power grids worldwide are not just protecting themselves against current threats but also preparing for a future where cyber-attacks are ever-evolving. This proactive approach is essential in today's digital age, where traditional security models no longer suffice, and adaptability is key to survival.

below, where zero trust will be applied to a segment rather than to the asset level. This is of course dependent on the type of assets and therefore the feasibility and potential maturity level of zero trust needs to be evaluated for the OT systems under consideration.

6.10 Vendor management

Vendor management can be regarded from different viewpoints. We rely on multiple vendors in most OT landscapes to provide equipment and services. Implementing a security evaluation in the procurement process is recommended to check whether the vendor meets or exceeds the security standards set by our organization. In addition, the NIS2 regulation might require the vendor to comply with the NIS2 when providing services to your organization. Creating a checklist or evaluation form will be very helpful. A prominent part is how a vendor will access its equipment when providing remote maintenance services.

A more strategic question is whether to rely on one vendor or on a multi-vendor landscape with regard to services and solutions. Vendors can have been compromised themselves, resulting in a supply chain attack like we have seen with Solarwinds. Using one vendor for a complete landscape will certainly be more efficient as everything integrates flawlessly, but when compromised, it can be difficult to value what still can be trusted and what is not. An independent monitoring solution or service would be wise to implement.

7. OT in different forms and sectors

In principle, OT is relatively the same across various sectors, as is the culture of the people working with OT. In most OT environments, bus network topologies are still very common, but the protocols used on them might vary from technology to sector. In an industrial environment, we will find protocols like Modbus and Profibus, while in vehicles, we'll have the Canbus, and in the maritime sector, NMEA.

When an organization plans to improve the security posture of its OT environment, it's good to understand the scope of OT. In industry, we assume the industrial control systems that run production, but do we also include the building management system (BMS) that handles the climate, the Safety Instrumented System (SIS), the elevators in the factory, and the BMS and elevators in the office building? In principle, all of the above.

And there's more to consider from a security perspective. The industry relies on IT and OT, but they also have laboratory environments in many cases. These labs can even be split up into quality assurance labs for production and labs for research and development, which contribute to the production of intellectual property.
We addressed the CIA triad for IT and OT SAI(C), but labs are kind of "in-between."
An R&D lab will have the operational characteristics of OT in terms of its assets and culture, but the IP produced is tied to the IT environment.

Another element to consider is product security. The automotive sector, shipyards, and manufacturers of medical devices are examples of organizations that rely on IT, OT, and labs and create products that themselves are considered OT.

8. The future is already there

As mentioned before, OT security is years behind IT security, while modernization from a business perspective is continuously ongoing. This means that organizations must not only act now with starting or scaling up their OT security improvement programs but also consider current developments when continuously evaluating the security controls to be implemented.

8.1 AI

Artificial intelligence (AI) has been around for some time, and machine learning has been used for years to monitor capabilities with automated responses. For example, endpoint security, firewalls, and advanced OT security monitoring. We do, however, need to recognize that adversaries can use the capabilities of AI for their attacks in various ways as well.

For AI to be successful, it needs access to data to learn, and as stated in 6.6, access is the key to securing the environment. A simple password is no longer adequate as it can be broken quickly. On the other hand, we must realize that OT still faces the challenges of achieving the highest availability, and to achieve that, immediate access to controls is a requirement. For this reason, we don't see much use of multi-factor authentication in OT, and in many cases, shared accounts are used and machines left unlocked for easy accessibility.
Improved access controls need to satisfy both security and operational requirements.

8.2 Digital twin

Digital twins help organizations to optimize efficiency and test improvements. From that objective, we see them moving into OT, but they have one big dependency to succeed. A twin can only be accurately built if we have complete insights into the environment, and as addressed in 5.8, we often don't in OT. The solution is to limit the scope of the digital twin to a smaller target for improvement.

An initiative to create a digital twin can benefit from an OT security improvement program as they share a common requirement: gaining insights into assets and their communication within the environment, like a production line. The costs for achieving these insights can be compensated by the results from efficiency-improving simulations of the digital twin.

But the relationship between a digital twin and security doesn't stop here. Next to the use for business perspectives, accurate digital twins can also be used for security purposes like, for example, penetration (PEN) testing. Pen testing is almost always a no-go in a life production environment as chances of disruption are high, resulting in downtime that is against the objective of maintaining high availability. Even pen testing during maintenance periods is tricky because testing, rollback, and validation need to be conducted quickly. Production lines generally do not have a complete test environment, so a physical twin is available. Pen testing in a digital twin

Digital twin proof of concept

At TCS I had the pleasure of realizing an innovative proof of concept from my vision regarding the use of digital twins for OT security perspectives.

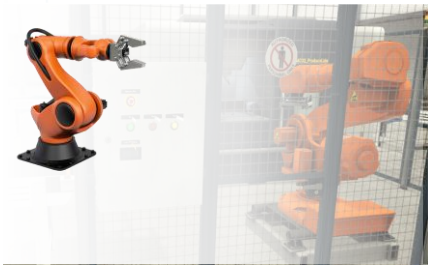

One part of the twin is built with physical components, like a robot arm, PLC and HMI and the other part is a graphically simulated 360 environment. Actions in the physical setup are mimicked in the virtual environment and vice versa. The virtual environment creates real OT network traffic that can control the physical components but can also be fed into OT security monitoring solutions that are available in the market. These monitoring solutions would identify assets in the virtual environment like they would in the physical world.

helps identify vulnerabilities that can then be examined and remediated with more focus on the maintenance period of the physical environment.

8.3 GenAI and Digital Twin

But what if we could generate a digital twin with Generative AI and virtual and augmented reality as used in the gaming industry? A digital copy of a factory where we could walk through and control the process using virtual HMIs.

This would be a powerful environment that can be used for many purposes, like training new operators, safety and security awareness training, cyber ranges, previously mentioned pen testing, and eventually, disaster recovery from the digital twin to the physical one.

8.4 Deception technology

Deception technology is available in the market and often comes in a “box” or virtual environment that pretends to be an OT environment, fooling potential adversaries. One might call it a honeypot on steroids, where the objective is to attract hackers, like a blue light attracts mosquitoes, with an environment that convinces them to be the jackpot. By closely monitoring this honeypot, the security solution quickly identifies attackers and can track and isolate them before they hit the real OT environment.

Combining the power of a digital twin and GenAI can increase the attractiveness and capabilities of such a honeypot, fooling even the best hackers.

9. About the author

Lucien Sikkens has over 24 years of experience in physical and cyber security in both IT and OT environments. He has worked in almost every sector and loves getting engaged and learning about an organization's business and processes. He uses this knowledge and experience in engagements and for tutoring clients, colleagues, and students.

Initially, Lucien studied electrical engineering, following in his father's and brothers' footsteps, but he switched to business administration along the way. His focus on the business perspective results in a humbler attitude towards security. It's business first, facilitated by automation and assured by security.

His passion for technology and automation and intrinsic motivation to improve an organization's resilience were his drivers to start a career in (cyber) security. He initially started in the financial sector but quickly broadened his perspective to other markets and focused on operational technology for the last 15 years. The attraction of OT is its versatility and security posture. While IT is generally the same for every organization, OT is used in many ways, feeding into his natural curiosity. Realizing that critical infrastructures and many other environments that our lives depend upon have become highly vulnerable to cyber-attacks is another reason for his dedication.

Lucien has a special affinity with the maritime sector and can often be found sailing at sea. OT, in general, is lacking in cyber security; it's even worse in the maritime sector, where cyber security awareness is almost zero. Practicing what you preach, he created a securely segmented network on his boat with secure remote access. His next paper will address the specifics of security in the maritime sector, including examples and potential innovations like a cyber threat radar for officers on the bridge.

Holding both CISSP and GICSP certifications, he has broad skills in securing both IT and OT environments. Over the years, he has implemented and run SOC operations, executed many assessments, created roadmaps, and provided advisory services to his clients. During his employment, he developed various managed security services and innovative concepts like the digital twin for OT from a security perspective and the CISO² concept, which provides CISOs with ad-hoc specialized knowledge and experience when needed on a subscription basis.

Since 2022, Lucien has held the role of Director of OT Security within the Global Consultancy Practice of Tata Consultancy Services (TCS). He engages with TCS customers worldwide and improves their OT security posture. If you would like to reach out, contact him at lucien.sikkens@tcs.com.

10. Resources

This document was written by the author based on his experience and knowledge, without the use of external sources. Included diagrams were created by the author and added photos were used from stock photo libraries or created with AI (Dall-E3) when confidentiality needed to be assured.

For readers that want to increase their knowledge of OT Security, I recommend the following sources;

- Book: Industrial Network Security by Eric D. Knapp and Joel Thomas Langill
  ISBN 978-0-12-420114-9
- YouTube/Internet: Excellent videos provided by Realpars.com explaining technology in the OT environment.