Secura
A BUREAU VERITAS COMPANY

REPORT

DISCOVERABILITY OF INDUSTRIAL SYSTEMS IN THE NETHERLANDS

By Arnau Ortega Gelabert, Mohamed Rezgui and Ralph Moonen, Secura | December 2021

Discoverability of Industrial Systems in the Netherlands

Introduction

With the emergence of COVID 19, 2020 and 2021 have been, without a doubt, extremely challenging years. The pandemic affected the way we use internet and has introduced new threats to both IT and OT systems due to a changed threat profile where remote work and remote access have become the norm due to travel restrictions and stay-at-home policies. As we become more and more connected, security becomes more challenging. Threat actors are also changing their methods and criminal business models, and seem to look towards OT and industrial systems for expansion of their criminal activities.

In the Netherlands, industrial systems and critical infrastructures are receiving more attention from regulators and CERTs. New legislation aims to bring more of that critical and industrial infrastructure under supervision of bodies such as Agentschap Telecom. The new NIS2 directive expands its scope and

Table of Contents

introduces liabilities to boards and directors of companies with regard to protecting assets from cyber attacks.

It is with this in mind that Secura set out to investigate the possibilities of finding weaknesses in internet-connected industrial systems. Such infrastructures rely on so-called Industrial Control Systems (ICS) and Supervisory Control And Data Acquisition (SCADA) networks. If an attacker obtains a way to disrupt or take over such systems, this could potentially have serious consequences across sectors resulting in financial loss, reputational damage, diminished consumer confidence, and even potentially threaten the safety of citizens and national security.

In this study, Secura devised and followed a methodology for identifying internet-connected cyber-physical systems (including industrial control systems and building management systems) and to gather data on potential vulnerabilities present.

This question was previously addressed in a 2019 study by the University of Twente (UT), titled "Online Discoverability and Vulnerabilities of ICS/SCADA Devices in the Netherlands" [UT2019]. This study attempted to answer several questions, including:

1. How many ICS/SCADA devices located in the Netherlands can be easily found by potential attackers?
2. How many of these devices are vulnerable to cyber attacks?

The study relied entirely on results from the Shodan search engine, and was therefore limited to what that search engine could present at that time. The second question was only answered in a very limited way in that study and this prompted Secura to try and improve those results.

Secura has extensive experience in performing security scans and data gathering projects, and decided to investigate these same questions in 2021, from a slightly broader perspective, and using scanning techniques to avoid reliance on 3rd party services such as Shodan. Secura's goals was partly to validate the 2019 UT study, and partly to satisfy curiosity about potential changes in the attack surface in the years since the UT study.

This document contains the results from Secura's investigation, and hopes to provide relevant insights to cybersecurity defenders, CERTs and the academic community as to what can be found and how to find them.

Management Summary

Industrial Control Systems (ICS) may have a lifespan of over 20 years. Therefore, they have traditionally been designed as independent systems, without sufficient security requirements, relying on air-gaps for security. Consequently, they are often not designed to deal with internet connectivity and current threats and targeted attacks. In order to analyse the status quo, it would be beneficial to know what the attack surface is for ICS in the Netherlands.

Our investigation focused on internet-connected ICS/SCADA systems in the Netherlands and attempted to answer the following questions:

- How many ICS systems are exposed on the public internet?
- How many of these systems are vulnerable or outdated?
- What information is exposed by the ICS systems discovered?
- How could the information available can be used to perform further attacks?
- What metrics can be used to evaluate evolution over time?

Using a combination of scanning and enumeration techniques – described in the methodology chapter – Secura gathered data that was then analysed in a variety of ways.

Secura
A BUREAU VERITAS COMPANY

Results

After scanning a number of protocols associated with (industrial) control systems, Secura found more than a thousand devices that could be classified as control systems in the Netherlands. The majority of these devices were building automation systems and seemed related to a small number of specific vendors. Since a specific protocol (Fox) also provides information on the names of the systems, it was possible to identify the owners of these systems. It was found that these are mostly schools and sports clubs, but also banks.

Other systems identified included pump stations and systems related to energy and solar panels. Since Secura did not perform any intrusive tests, information about potential vulnerabilities was limited to gathering version information and studying easily verifiable issues such as SSL/TLS issues and web interface issues (missing headers, administrator login pages). Approximately twenty percent of the identified devices indeed also offered a web interface, and virtually all of these web interfaces seemed to be vulnerable to one or more issues that Secura checked for, confirming that these types of devices are often outdated or not well-protected.

Secura's research found more devices than Shodan.io, and more devices than [UT2019]. There are several logical explanations for this however, (most notably the fact that internet connectivity has increased in the past two years) and the difference is not large enough to invalidate results of [UT2019].

Given that the attack surface of industrial systems is constantly evolving, it is interesting to track how this happens over time. Secura therefore intends to repeat this research periodically and has recommendations for tracking this on a national level.

Risks

The risks of having ICS equipment connected to the internet and discoverable for everyone, is difficult to quantify. However, it is generally bad practice to expose industrial protocols, version information, configuration information and administrator login pages to the internet. If an attacker would find that a targeted company or sector relies on specific vendors and devices, they might choose to point their attention to those devices and identify vulnerabilities in them. Secura identified several cases where this seemed to be a relevant risk, mainly, again, related to building automation systems.

The vulnerabilities that Secura identified in specific equipment carry certain well-known risks. SSL/TLS vulnerabilities can allow an attacker to listen in on sensitive communication, while certain web-based vulnerabilities might allow attackers to take over administrative functions. If such attacks are performed on a large enough scale, the risk of disruption exists.

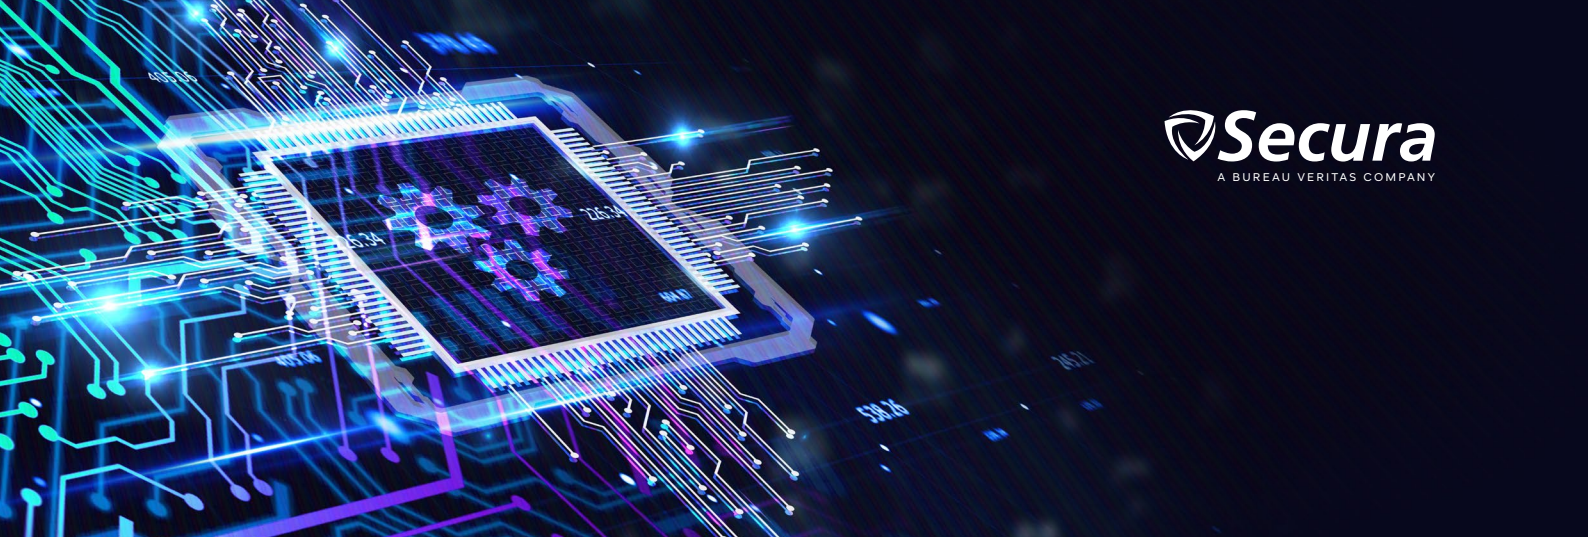

Conclusions

Secura set out to answer several research questions. In summary, the following answers can be formulated:

1. How many ICS systems are exposed on the public internet?

It is possible to identify these systems by the protocols and services they publicly offer to the internet, sometimes without sufficient authentication. At least one thousand of such devices were found.

2. How many of these systems are vulnerable or outdated?

Virtually all of the identified systems showed one or more vulnerabilities, and many of those vulnerabilities were related to the use of outdated protocols.

3. What information is exposed by the ICS systems discovered?

Systems often expose version information. In the case of specific protocols, they also expose the exact location of the device including name and function.

4. How can the information exposed be used to perform further attacks?

Attackers could use this information to target specific assets, or become motivated to research and identify specific types and vendors of devices that look vulnerable.

5. What metrics can be used to evaluate evolution over time?

Absolute numbers of identifiable ICS systems are a prime metric. Secondary metrics could include the number and risk rating of easily identifiable vulnerabilities.

Recommendations

Given that Secura's research showed a few potentially sensitive targets, it must be assumed that an attacker with sufficient motivation will also be aware of these. Finding such targets and mitigating the risks to them (which would start with identifying and notifying vulnerable asset owners in most cases) is a task that currently is not formally executed on a national level by a governing body. The discussion around this topic currently looks towards the new NIS2 directive, where the scope has been extended to also include important organizations (and not just essential organisations) and no longer limits its reach to only specific sectors. Soon, all important organisations will potentially fall under the regulation.

The new directive mandates implementation of proper security measures, monitoring and reporting. It also enables better sharing of information between stakeholders. These are important improvements and will certainly benefit cyber resilience, including for industrial systems. However, without a system of automatic vulnerability and attack surface mapping of those assets, cyber resilience will remain predominantly reactive in nature. It is for this reason that Secura recommends implementing a permanent, continuous, preventive scanning system on a national scale. Such a system would try to identify vulnerable assets (using similar methodologies as described in this study) that can be associated with organisations that are important or essential to the Dutch economy and state (and by extension, fall under the NIS2 directive). Secura also recommends that this is aligned with developments in other EU-countries and with sectors that might be underexposed to NIS2 such as software companies that carry their related supply-chain risks.

Methodology

Scanning a country's IP space is always a question of balance between thoroughness and speed, while respecting ethical and legal boundaries. Although there is no technical limitation that prevents scanning all TCP and UDP ports exposed to the internet, the time this takes would render the task unfeasible. To overcome this limitation, this study designed a series of steps to progressively narrow down the amount of IP addresses to scan. The following sections describes the actions taken to implement this approach.

Scope Selection and Initial Scan

Theoretically, the IPv4 space consists of 4.2 billion addresses. However, not all of these IP addresses can be used to identify individual hosts. Certain IP addresses have a special purpose, like private or broadcast addresses. When they are removed, the number goes down to 3.7 billion.

Furthermore, each country has been assigned specific ranges of IP addresses over the years by the Internet Assigned Number Authority (IANA). For the case of the Netherlands, a number of 51.6 million IPv4 addresses are officially assigned.

This study focused only on the Dutch IP space. To define this scope, the list of IP addresses assigned to the Netherlands was obtained from Ipinfo[1], a service that maintains a public registry of the IPs ownership. According to previous experience within Secura, it is known that certain Internet Service Providers within the Netherlands are sensitive to wide port-scans on their (mainly residential) customers' ranges. To limit the reach of our study regarding these residential IP ranges the following strategy was implemented. First, a list of all Autonomous Systems (AS) belonging to known ISPs was gathered. Then, information about all these AS was retrieved using *whois* queries.

Finally, any AS record that returned keywords like "Customers" or "ADSL" was removed from the list of AS systems to scan. We certainly realize that this also might impact the results but the assumption here is that there will not be a huge number of industrial assets on residential and civilian IP ranges anyway. However, we did include certain mobile ranges due to the assumption that industrial assets could very well be connected using M2M networks using 2G/3G/4G networks (and this was indeed the case).

The final list of networks obtained after applying the mentioned operations was passed to *masscan*[2] in order to perform a first targeted portscan related to ICS and web services. The final list of ports can be consulted Appendix A.

ICS Information Gathering

Masscan only returns what ports are open on each hosts and not which protocol is used, and due to its nature, high numbers of false positives were expected. Zgrab2[3] was used to reduce the number of such events by scanning only the top-5 protocols in terms of prevalence, and perform a first interaction with the hosts. With this, we managed to discard a large number of presumed false-positives and get basic information about the device like: model, version number, operating system and in some cases the owner. The list of ICS protocols that were queried using Zgrab2 can be consulted in Appendix B.

Web Crawling and SSL Assessment

Web-services were also considered relevant and worth to scan, enumerate and identify during this study. The reason for that, is while it is common to have incoming connections to non-standard ports blocked on edge firewalls, web ports are often allowed. At the same time, ICS appliances and related equipment sometimes expose a web portal for management together with the ICS port. These occurrences potentially open the door to unintended access to sensitive equipment for anyone on the internet whether through publicly available vulnerabilities or even social engineering (phishing) or brute force attacks on the web interfaces.

[1] https://ipinfo.io/countries/nl
[2] https://github.com/robertdavidgraham/masscan
[3] https://github.com/zmap/zgrab2

In order to identify potential candidates and reduce the number of hosts to query, nmap[4] scripts were used to first query http-title and server headers of all hosts where an active ICS port was discovered.

The next step consisted of filtering the results using a custom wordlist that contained ICS related terms both in Dutch and English. The list also contained vendor names for some products.

The results obtained through this enumeration attempt were considered to have a high probability to be related to ICS systems. Therefore, more interaction was performed with those hosts.

First, a screenshot was obtained for each web application found from the previous list using Aquatone[5]. Next, Testssl[6] script was used to identify vulnerabilities related to the encryption layer used (TLS/SSL). Finally, common HTTP misconfigurations such as missing headers and security attributes were scanned for.

VNC Protocol

The study also looked at the VNC protocol. This protocol is sometimes implemented to provide a graphic interface to control a device remotely. Due to its simplicity and lightweight implementation it is a protocol that has been used in the past to manage ICS. Moreover, legacy equipment sometimes does not implement any authentication and anyone could connect and interact with the device. We therefore also screenshotted any unauthenticated VNC server, similar to what Shodan.io also does.

SNMP Protocol

The SNMP protocol is used to provide network management services between a central management console and network devices such as routers, printers, and also ICS devices such as HMIs (Human Machine Interfaces) and PLCs (Programmable Logic Controllers). Although SNMP is an extremely useful service for maintaining a network, it is often very weakly secured. Both versions 1 and 2 of SNMP make use of unencrypted passwords such as the default community string "public" to both read and sometimes also to configure devices (including devices such as PLCs). Devices offering public SNMP can often easily be fingerprinted because they advertise a system identification string.

Data Analysis

All results obtained were imported into an Elastic database. This allowed us to correlate and aggregate the results and obtain meaningful insights into the data collected.

Additionally, the hosts that exposed vulnerabilities were manually classified into sectors. Example sectors are, education, public government, critical infrastructure among others. This classification was based on identifiable information gathered from the hosts itself or by querying public information databases.

[4] https://nmap.org/
[5] https://github.com/michenriksen/aquatone
[6] https://github.com/drwetter/testssl.sh

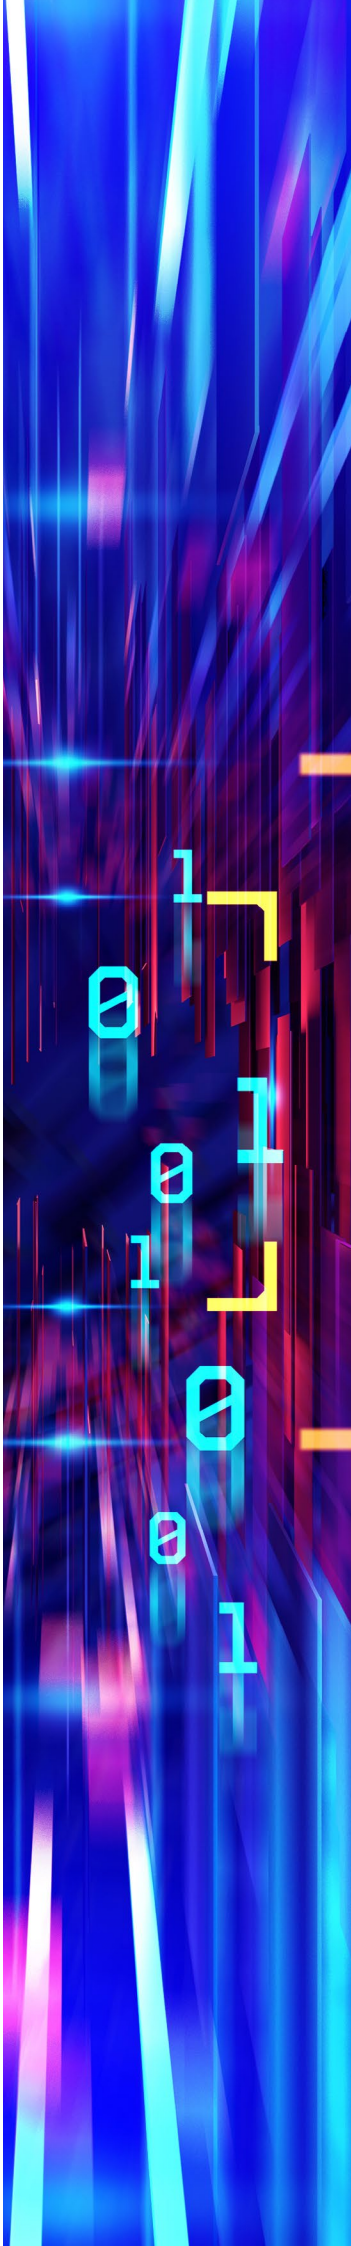

Results

Masscan

The results in the following table show the open ports returned by *masscan*. Due to the nature of this type of scanning, a high number of false positives were expected. Nevertheless, the results were in the same order of magnitude as in previous experiences within Secura and also in line with [UT2019] although slightly higher. A likely explanation for this could be that since 2019, more devices have become internet-connected.

As can be seen, the most common ports belonged to web services, by a considerable margin. The first non-web port was DNP3 in position 6 for TCP/20000, followed by VNC in position 9. Please note that an open port does not mean that it necessarily also uses the protocol associated with that port. In this phase of the research we did not interact to validate the actual protocol yet.

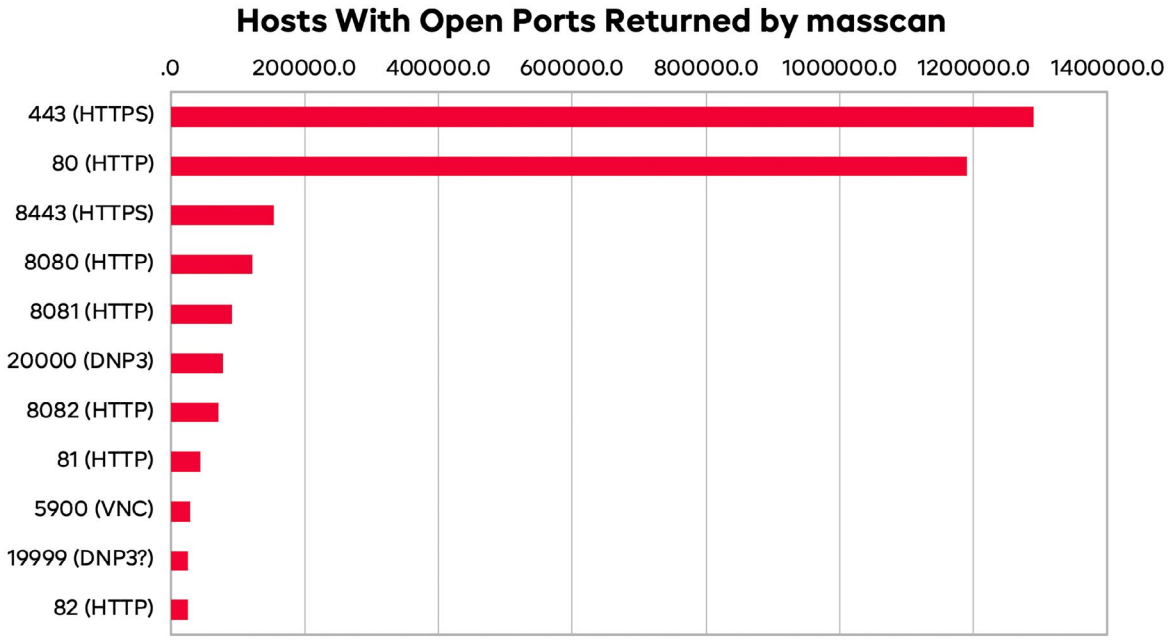

Figure 1: Absolute numbers of open ports scanned

ICS (with zgrab2)

As explained above, zgrab2 and a limited interaction (banner grabbing or simple readout) was used to gather as much information as possible from a selection of specific ICS protocols. The list of hosts scanned was based on *masscan* results.

In the graph below, the number of hosts that returned a successful handshake with meaningful information is shown. The Fox protocol is the most common one, followed by Siemens S7 and Modbus. These results are again in line with previous public research. The positive presence of ICS related information confirmed these hosts were true ICS devices.

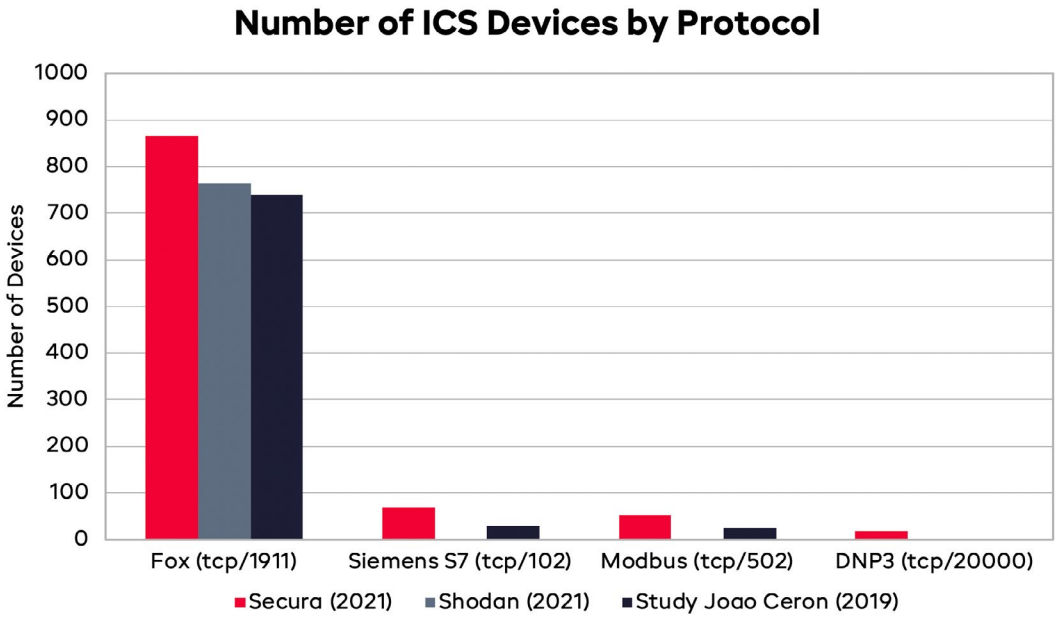

Figure 2: Identified ICS protocols compared with other sources

The following shows an example of the data collected by zgrab2 for the Fox protocol. This protocol was the most verbose one and allowed to gather considerable information related to the version installed and the device name assigned by the owner. Later on, this information was used to determine the sector for its device.

```
"station_name": "<redacted>",
"host_address": "192.168.1.132",
"sys_info": "bog 61[<bog version=¥"1.0¥">",
"app_version": "3.7.108.2",
"os_version": "6.4.1",
"language": "nl",
"time_zone": "Europe/Amsterdam",
"version": "1.0.1",
"host_id": "Qnx-NPM6-0000-1471-2949",
"brand_id": "<redacted>",
"app_name": "Station",
"hostname": "192.168.1.132",
"vm_version": "1.5.0_81-b02",
"is_fox": true,
"os_name": "QNX",
"id": 1535,
"vm_name": "Java HotSpot(TM) Client VM"
```

These devices advertise this information publicly, and as can be seen it contains information that could be of interest to an attacker. The *station_name* provided a unique identifying name of the target in all cases.

Web Interface Assessment (SSL/TLS Analysis)

One of the main assumptions of this study was that ICS devices might expose a web-based management interface together with the ICS protocol communication port. The following graph shows which ones of the ICS hosts queried by zgrab2 also exposed a web server.

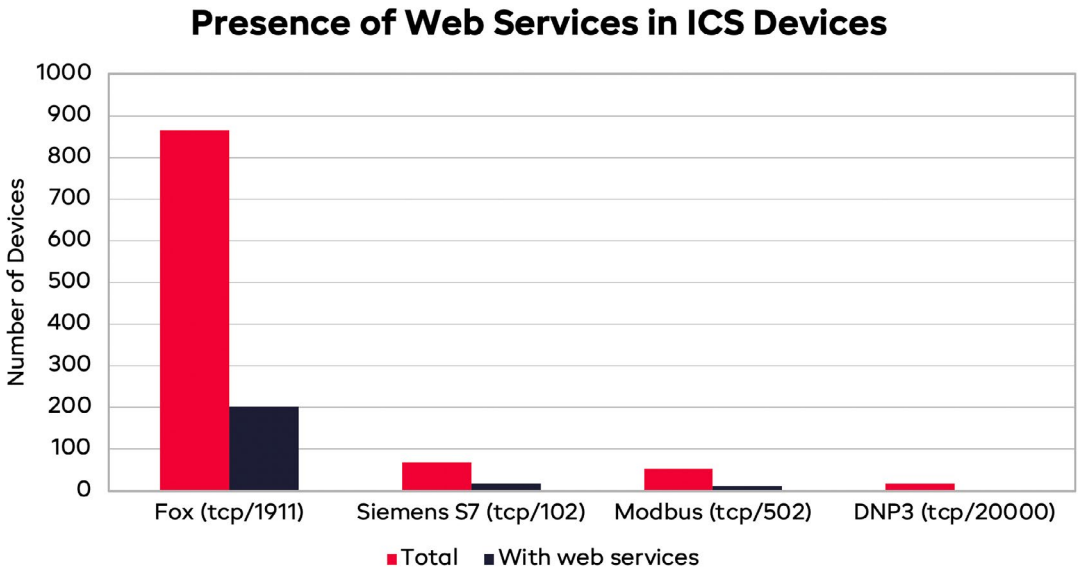

Figure 3: prevalence of web interfaces per common ICS protocol

However, an ICS protocol can be implemented by multiple vendors. With the objective to better understand and compare across brands the *brand_id* was used to identify the device vendor.

As can be seen, the presence of HTTP transport for the web management interfaces was present across all identified ICS manufacturers. However, its presence was more predominant for certain brands. This can be explained by the different nature of the device. For example, some devices can have a web service enabled by default while others do not offer such option.

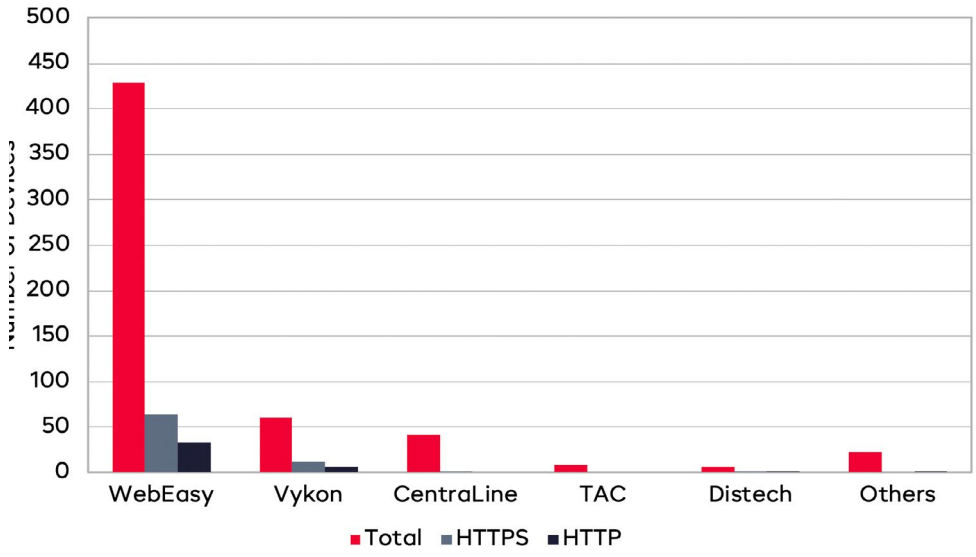

Figure 4: HTTP-enabled equipment brands compared

Interestingly, the majority of these brands focus on building automation. Given that building automation security has received comparatively less attention in literature and publications, this is a hint that building automation is lagging in security hygiene more than other parts of the industrial market.

Devices with a high probability of being an ICS were scanned using the *testssl.sh*[7] script on ports 443 and 8443. In the graph below it can be seen how the lack of security headers is the most common vulnerability followed by the support of outdated SSL/TLS versions.

The results also reconfirm the fact that many ICS systems are updated often since many of the identified issues and TLS/SSL deficiencies have been very well described and widely known for more than a decade.

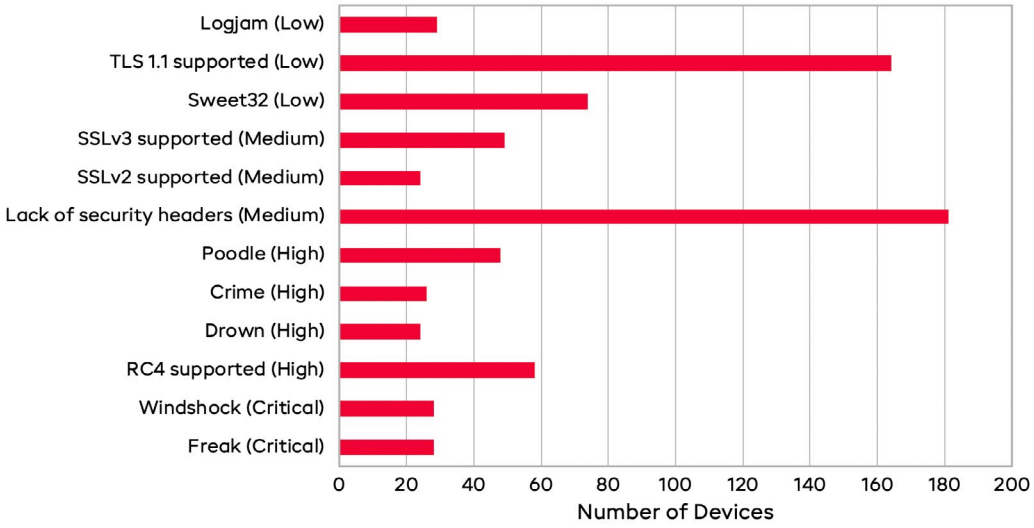

Figure 5: Vulnerabilities related to web services on ICS devices

This information can be represented by vendor. As can be seen, all of them expose the same distribution of vulnerabilities in relation to its presence. One explanation of this fact could be these vulnerabilities were discovered after the device was introduced. Then, the lack of support by the vendor or the lack of maintenance by the owner of the device rendered the device vulnerable. The presence of critical and high-risk vulnerabilities across a number of potentially important devices and the fact that these are exposed to the public internet, could be a reason for concern.

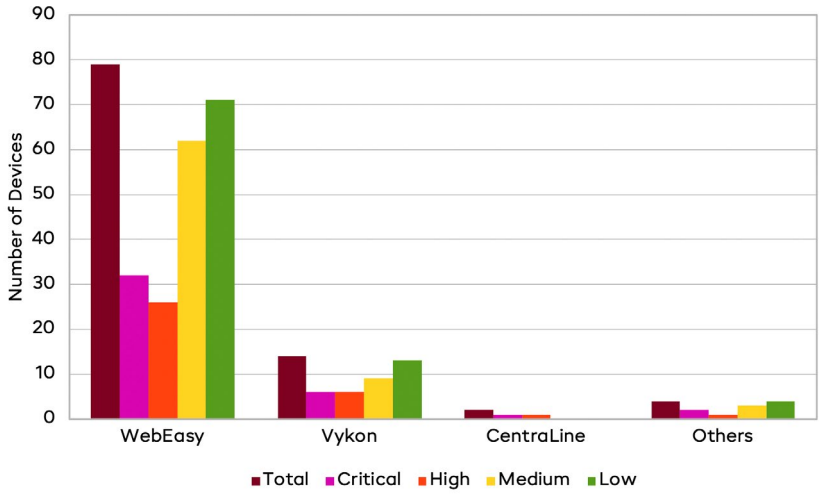

Figure 6: Criticality of web-based vulnerabilities per vendor

[7] https://testssl.sh/

Cases by Sector

It was possible to identify several systems through the usage of the Fox protocol. The following graph shows a list of the sectors that were using vulnerable products from the vendors mentioned above:

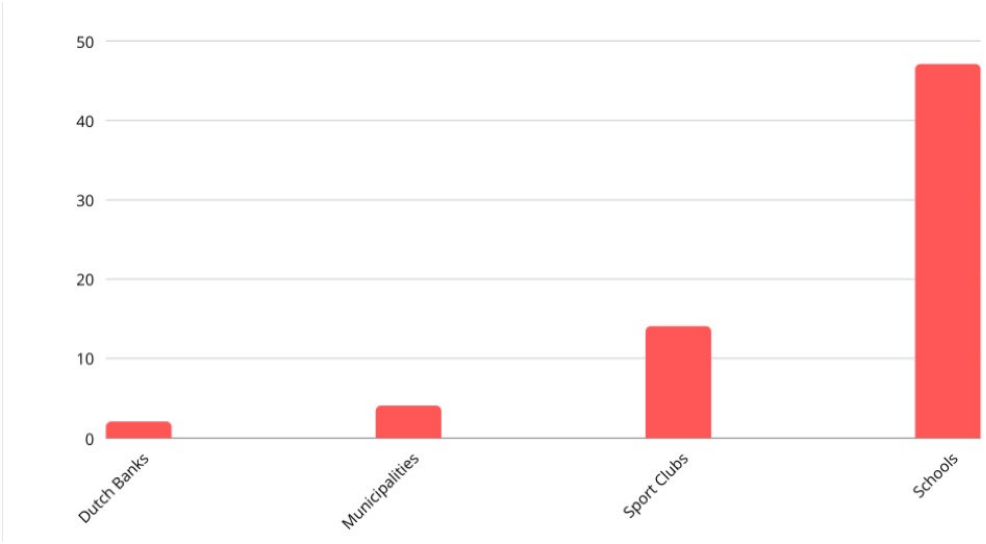

Figure 7: Usage of Fox protocol related to sector

Some Examples of Discoverable Assets

Obviously it is possible for anyone to repeat our methodology, and while we do not want to endanger any asset owners, we do feel it can be illustrative to include several anonymised examples of identified assets.

While we have not attempted to log into or exploit any of these systems due to legal and ethical limitations, we note that an attacker might be less inclined to follow such standards.

```
└─$ grep "HMI" /home/mohamed/Desktop/20210211-OT-threat-exposure/02 Workfolder/snmp_results.txt
89.[redacted] [public] Siemens, SIMATIC HMI, TP700 Comfort, 6AV2 124-0GC01-0AX0, HW: 0, SW: V 13 0 1
87.[redacted] [public] Siemens, SIMATIC HMI, TP1200 Comfort, 6AV2 124-0MC01-0AX0, HW: 0, SW: V 14 0 1
91.[redacted] [public] Siemens, SIMATIC HMI, TP900 Comfort, 6AV2 124-0JC01-0AX0, HW: 0, SW: V 14 0 1
```

Figure 8: HMIs running SNMP service with default community string

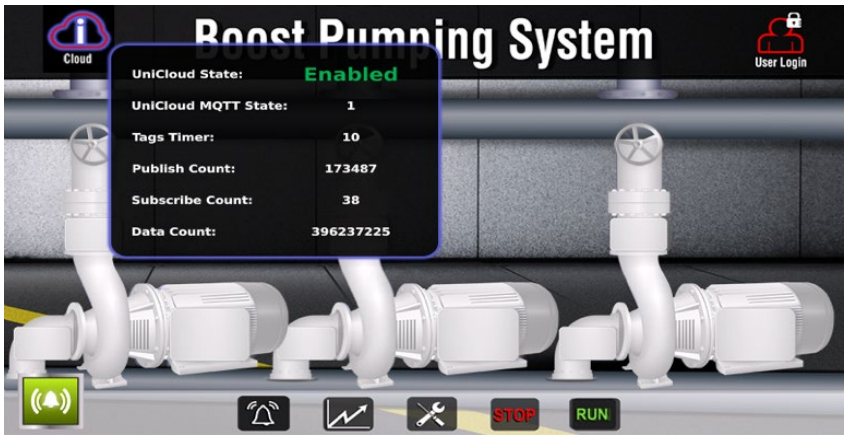

Figure 9: A pump system

An example of a pump system that provided a VNC connection of their console. Please note that certain systems could be test setups, development environments or honeypots, and since no interaction took place we were not able to ascertain the production status of any results.

ip	[redacted]
masscan.ports	80, 1,911
zgrab2.protocol	fox
zgrab2.result.app_name	Station
zgrab2.result.app_version	3.8.401
zgrab2.result.auth_agent_type	fox
zgrab2.result.brand_id	webeasy.products
zgrab2.result.host_address	192.168.1.128
zgrab2.result.host_id	Qnx-NPM6E-0000-16D4-F3F6
zgrab2.result.hostname	192.168.1.128
zgrab2.result.id	131
zgrab2.result.is_fox	true
zgrab2.result.language	nl
zgrab2.result.os_name	QNX
zgrab2.result.os_version	6.5.0
zgrab2.result.station_name	[redacted]
zgrab2.result.sys_info	bog 61[<bog version="1.0">
zgrab2.result.time_zone	Europe/Berlin
zgrab2.result.version	1.0.1
zgrab2.result.vm_name	Java HotSpot(TM) Embedded Client VM
zgrab2.result.vm_uuid	11eb5cad-0378-d030-0000-000000009485
zgrab2.result.vm_version	25.161-b01
zgrab2.status	success
zgrab2.timestamp	Feb 16, 2021 @ 18:05:36.000

Figure 10: Building control system of a Dutch bank

Conclusion

Secura has used the results from the scans to attempt to answer the research questions. The following can be concluded (all within the context of The Netherlands):

How many ICS systems are exposed on the public Internet?

According to the information gathered, more than one thousand confirmed ICS devices were exposed to the public internet. This number comes from the interaction with specific protocols and ports. In fact, the researchers think this number could be considerably larger. The reason for this comes from the definition of ICS. As the name implies, a given device is considered ICS when its purpose is related to the control or management of industrial equipment. While this is clear when the service exposes certain ICS specific protocols, a management portal exposing industrial processes, readings or controls should also be considered part of this group. However, the latter imposes serious limitations when it comes to easily classifying its purpose. Devices that do not allow themselves to be classified easily have not been included in our results. This means that the actual number is probably much higher.

How many ICS systems exposed on the public internet are vulnerable or outdated?

The scans reported a considerable number of vulnerabilities across the different vendors and protocols.

The uniformity in the distribution of vulnerabilities across the vendors and protocols indicates that no particular product has been proven less vulnerable than others. One possible explanation for such fact is that vulnerabilities were discovered after the product was launched. Then, the lack of software updates by part of the vendor, or the lack of proper maintenance by failing to keep the system up to date by the owner, rendered the device vulnerable. Additionally, the assumption exists that many devices use resource-constrained computing systems and were never designed to be internet-connected.

What information is exposed by the ICS systems discovered?

The type of information exposed by ICS device was diverse. However, most of the time version information, device type, and data related to a particular industrial process was exposed. In some cases operational values of certain machinery could be seen and possibly altered. In other cases, only connection information could be modified which would render the device inaccessible creating potential downtime. Moreover, many building management devices exposed enough information to identify the specific building.

How can the exposed information be used to perform further attacks?

The unauthorized access to ICS devices could range from inconsequential to an action resulting in a critical impact in the business continuity. However, in most cases access read-only or behind an authentication page. In order to cause impact to processes or property, an attacker would need to use a non-public vulnerability to get full access to a device or an internal network.

Yet, the fact that such kind of devices are reachable from the public internet should be considered a vulnerability. Weaknesses in the software or firmware could have a large impact. Depending on the target, an attacker might become sufficiently motivated to go out and research such vulnerabilities. The (growing) presence of certain ICS devices in parts of the production chain could therefore awaken the interest of threat actors who would then face an exposed ICS landscape with an outdated technology stack ready to be compromised.

What metrics can be used to evaluate evolution over time?

This study suggests several metrics that could be used to track the security posture evolution of internet-connected ICS systems. For example, taken into account that it is not advised to expose ICS devices to the public internet, their mere presence could be used as a metric to evaluate the maturity of ICS security in the Netherlands. However this must be corrected for the dominance of a specific vendor if that is identified.
Furthermore, as mentioned before, several ICS devices expose a web interface for management. The use of old protocols for the encryption layer could give a hint of a device software age which could be used as a metric to track updating.

The value of these metrics would be highest however if they can be measured over a longer period of time. This is also a good reason to suggest a periodic and frequent reproduction of this study.

Questions?

Do you have any questions about ICS security or do you have a question about our other services?

Contact us at info@secura.com or call us at +31 (0)88 888 3100.

Appendix A. Ports Scanned

The following table lists the ports that were scanned for this research paper. It is a selection of ports that are associated with industrial systems. In future research we intend to expand this list to include other relevant protocols.

Portnumber	Protocol
TCP/80, TCP/443, TCP/8443, TCP/8080	HTTP
Alternative HTTP: 81,82,8081,8082 etc.	HTTP
TCP/80, TCP/443	OPC UA XML
TCP/102	ICCP
TCP/102	IEC-61850
TCP/102	Siemens S7Comm
UDP/161	SNMP
TCP/502	Modbus TCP
TCP/1089-1091, UDP/1089-1091	Fieldbus HSE
UDP/1628, UDP/1629	LonTalk
TCP/1911	Fox (Tridium/Niagara)
UDP/2222, TCP/44818, UDP/44818	Ethernet/IP
TCP/4059, UDP4059	DLMS/COSEM
TCP/4712, UDP/4713	IEEE C37.118
TCP/4840	OPC UA
TCP/5094, UDP/5094	HART-IP
TCP/5900	VNC
TCP/20000, UDP/20000	DNP3
TCP/34962-34964, UDP/34962-34964	PROFINET
UDP/34980	EtherCAT
UDP/47808	BACnet/IP
UDP/55000 – 55003	FL-net

Appendix B. ICS Ports Scanned by Zgrab2

- Fox, 1911
- Siemens S7Comm, 102
- Modbus, 502
- DNP3, 20000

Secura
A BUREAU VERITAS COMPANY

BUREAU VERITAS
Shaping a World of Trust

Karspeldreef 8
1101 CJ Amsterdam
Netherlands

Vestdijk 59
5611 CA Eindhoven
Netherlands

T +31 (0)88 888 3100
E info@secura.com
W secura.com

Follow us on