FORESCOUT RESEARCH | VEDERE LABS

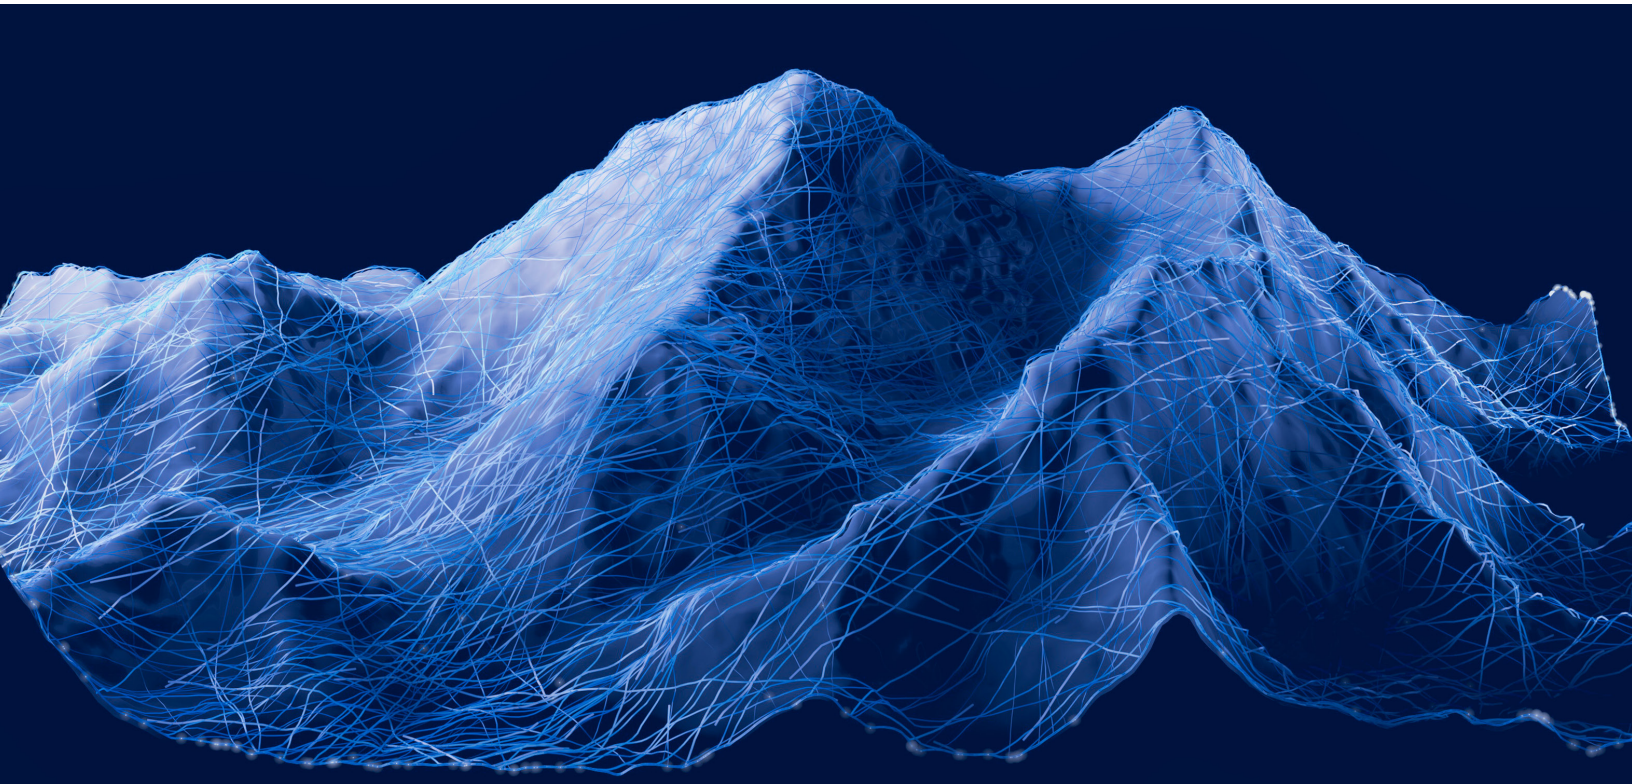

CONCLUDING OT:ICEFALL

New vulnerabilities and a retrospect on OT security design and patching

Jos Wetzels

Abdelrahman Hassanien

June 20, 2023

Table of Contents

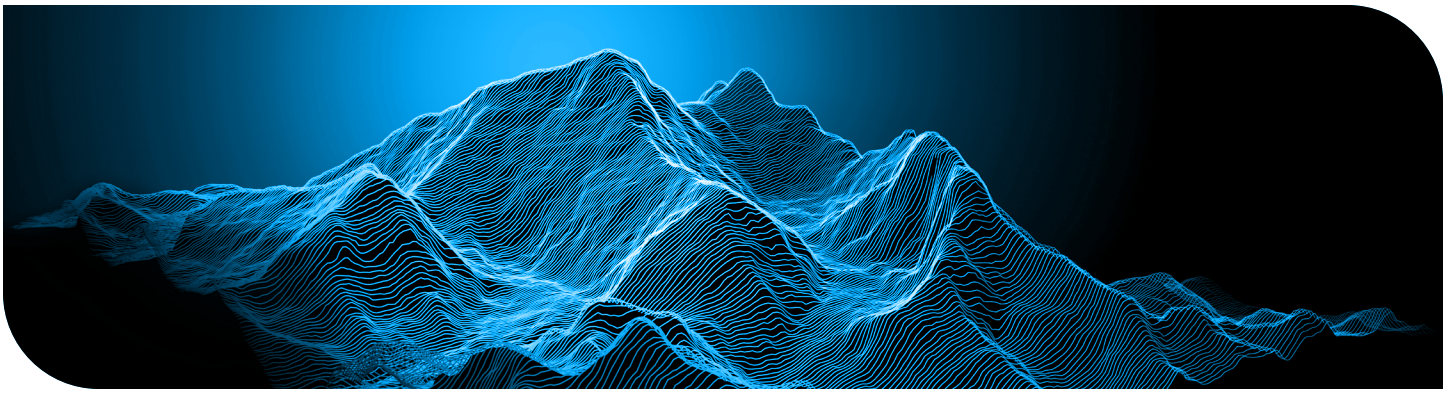

Executive summary

Since 2020, Forescout Vedere Labs has been finding and disclosing high-impact vulnerabilities in IT, OT and IoT devices. We started with **Project Memoria**, a large analysis of TCP/IP stacks. In 2022, we released **OT:ICEFALL**, which focused on insecurity-by-design in OT and included 61 vulnerabilities, which affected more than 100 individual products manufactured by 13 vendors.

Along the way, we have twice shared lessons learned, **first** focusing on the dangers of silent patching, the predictability of vulnerabilities and the unresponsiveness of vendors; and **later** focusing on the long-term positive effects of vulnerability research, such as the improvements of standards and the use of better test cases for development.

In this report, we present the three last vulnerabilities in OT:ICEFALL and conclude the project after one year by presenting a new retrospect.

New vulnerabilities:

- CVE-2022-46680 is the last issue found in the original OT:ICEFALL research and not made public at the request of the vendor. This vulnerability concerns the plaintext transmission of credentials in a protocol used by power meters from Schneider Electric. By abusing this vulnerability, attackers can compromise credentials and gain control of vulnerable devices.
- CVE-2023-1619 and CVE-2023-1620 are new findings on WAGO controllers using the CODESYS runtime. These vulnerabilities allow for denials of service (DoSs) on the affected devices either by sending specific malformed packets or by sending requests after a user has logged out.

In the retrospect, we focus on three new insights we can distill from the year-long research project and disclosure of these vulnerabilities in OT products:

- **Vendors still lack a fundamental understanding of secure-by-design.** Apart from typical implementation flaws such as plaintext and hardcoded credentials or key material, we encountered several recurring design issues that demonstrated how many OT vendors lack a fundamental understanding of basic security control design principles.
- **Vendors often release low-quality patches.** Incomplete patches for several issues have led to new vulnerabilities being discovered, which exemplifies how a bad patch *increases* risk instead of decreasing it.
- **Vendors must improve their security testing procedures.** The shallow nature of many vulnerabilities found casts doubt on the quality of the testing these products are supposed to undergo. In addition, some vendors have a certified software development lifecycle, which leads us to wonder how the bugs were missed by those vendors in the first place.

Following are details about the newly disclosed vulnerabilities and our research takeaways, including the state of OT product security, the vulnerability management implications for asset owners and why we recommend a consequence-driven (vs. likelihood-driven) risk approach to OT security.

1. New vulnerabilities

Table 1 summarizes the new vulnerabilities we are disclosing. CVE-2022-46680 is the last issue found in the original OT:ICEFALL research and not made public at the request of the vendor. CVE-2023-1619 and CVE-2023-1620 are new findings on WAGO controllers using the popular CODESYS V2 runtime.

Table 1– New vulnerabilities

CVE ID	AFFECTED DEVICES	DESCRIPTION	CVSS V3.1	POTENTIAL IMPACT
CVE-2022-46680	Schneider Electric ION and PowerLogic power meters	The ION/TCP protocol transmits a user ID and password in plaintext with every message, allowing an attacker with passive interception capabilities to obtain these credentials.	8.8	Compromise of credentials
CVE-2023-1619	WAGO 750 controllers	An authenticated attacker could send a malformed packet to trigger a device crash.	4.9	DoS
CVE-2023-1620	WAGO 750 controllers	An authenticated attacker could send packets in a specific sequence to trigger a device crash.	4.9	DoS

ION and PowerLogic power meters provide power and energy monitoring. According to **Schneider Electric**, these products are deployed in sectors such as manufacturing, energy, water and wastewater systems. **WAGO 750** is a line of automation controllers with variants supporting several different protocols, such as Modbus, KNX, Ethernet/IP, PROFIBUS, CANopen, BACnet/IP, DeviceNet and LonWorks. The controllers are **known to be used** in sectors such as commercial facilities, manufacturing, energy and transportation.

1.1. CVE-2022-46680

The Schneider Electric ION and PowerLogic product lines use the ION/TCP protocol on port 7700/TCP for communications between a master terminal and energy monitors. This protocol transmits a user ID and password in plaintext with every message, provided this feature is enabled, allowing an attacker with passive interception capabilities to obtain these credentials.

An attacker who obtains ION or PowerLogic credentials can authenticate to the ION/TCP engineering interface as well as SSH and HTTP interfaces to change energy monitor configuration settings and potentially modify firmware. If the credentials in question are (re)used for other applications, their compromise could potentially facilitate lateral movement.

Some PowerLogic ION meters have functionality called **ModemGate** and **EtherGate**, which allows the use of the meters' internal modems to transmit data from serial devices (RS-232 or RS-485) and Ethernet networks, respectively, to nested (third-party) device networks behind the meter. These features, especially combined with CVE-2022-46680 in case authentication is enabled on nested devices, could potentially lead to a **deep lateral movement** scenario.

To address CVE-2022-46680, as of firmware version 4.0.0 (ION9000, PM8000, ION7400), Schneider Electric introduced **Secure ION** (on port 7443/TCP), supported by the latest version of ION Setup and in future releases of Power Monitoring Expert (PME). Secure ION uses a TLS tunnel between the energy monitor and master terminal, but is *disabled by default.* In addition, ION sessions (both regular and Secure ION) now use a session token ID to authenticate ION requests to a device.

Devices that support Secure ION use **self-signed certificates** by default and require CA-signed certificates to be uploaded during commissioning. Invalid security certificates will trigger a warning on part of the engineering software. After a connection has been established, the user sends credentials to the device, which returns a unique, non-transferable token to the engineering software. Subsequent ION requests are validated using the token ID (with session timeout of 5 minutes) instead of the credentials.

While we have not looked into the new session mechanism in detail yet, and the approach is not secure-by-default, it does at least present a shift towards secure-by-design for the ION/PowerLogic product lines and attempts to address the root cause of the original vulnerability.

1.2. CVE-2023-1619 and CVE-2023-1620

CVE-2023-1619 exemplifies the poor design of the protocol parsers integrated in the WAGO 750 controllers and allows authenticated attackers to crash the device by sending malformed packets. There have been several similar bugs with these protocol parsers in the past. **CVE-2019-19789**, for example, led to denial of service due to a null pointer dereference in products using CODESYS V2 runtime. Other examples include **a set of issues identified in 2019** that could lead to remote code execution and **CVE-2020-12522**, which was discovered in 2020 but fixed by a silent patch in 2017.

CVE-2023-1620 is an example of insufficient session expiration (CWE-613), where an attacker can crash an affected device by sending specific requests after being logged out of the device. This bug could be fixed by the controller closing the connection after logout to avoid memory-based vulnerabilities.

After triggering any of the vulnerabilities, the affected device must be manually rebooted to return to its operating state.

We found the vulnerabilities on a **WAGO 750-862** controller running firmware version FW0750 0862 V010504, released on February 10, 2022. However, both vulnerabilities affect WAGO's use of the CODESYS V2 runtime, which means they affect more devices than we originally tested as they stem from the flawed integration of a supply chain component.

1.3. Distribution of devices

Although these devices are typically not supposed to be exposed online, we see between 2,000 and 4,000 potentially unique devices directly accessible when querying the **Shodan** search engine, as shown in Table 2. The most popular exposed protocol is HTTP for WAGO controllers and Telnet for ION meters. WAGO controllers are most popular in Europe (especially Germany, Turkey and France), while ION meters are most popular in North America.

Table 2 – Devices exposed on Shodan. Query on April 25, 2023

QUERY	DEVICES	TOP 5 COUNTRIES
"Description: WAGO 750" (SNMP)	57	Germany: 22 Poland: 9 Italy: 7 Switzerland: 6 France: 3
"Friendly Name: WAGO" (KNX)	23	Hungary: 6 Belgium: 5 Italy: 4 Germany: 2 Poland: 2
"Product name: WAGO 750" (EtherNet/IP)	44	Croatia: 16 Italy: 12 Germany: 6 Romania: 3 Turkey: 3
"Server: WAGO_Webs" (HTTP)	1,428	France: 299 Germany: 148 Poland: 117 Japan: 116 Spain: 110
http.html:"WAGO ethernet" (HTTP)	2,297	Turkey: 728 Germany: 577 Italy: 182 France: 145 Japan: 108
port:23 "ion" (Telnet)	33	Mexico: 14 US: 7 Canada: 4 Ecuador: 4 Argentina: 2
Powerlogic FTP (FTP)	5	US: 4 Germany: 1
http.html_hash:2039835388 (HTTP)	11	US: 11 Mexico: 1

On the Forescout Device Cloud – a repository of data from 19 million devices monitored by Forescout appliances – we see around 500 WAGO controllers and 500 ION power meters, as shown in Figure 1. Both types of devices are most commonly seen in manufacturing, but they are also popular in utilities and healthcare, in the latter case mainly for building automation.

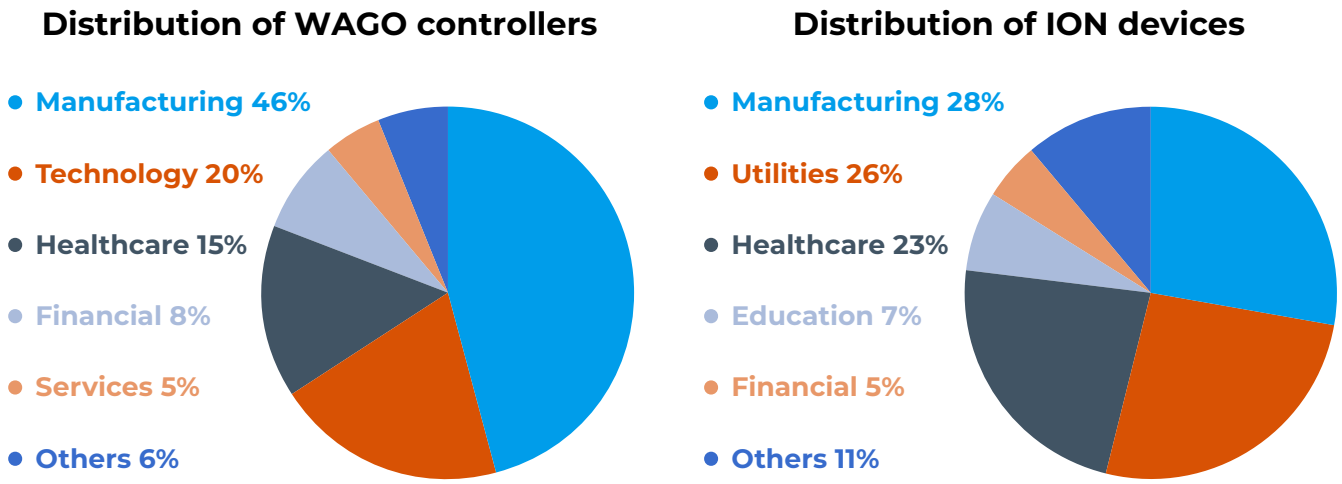

Figure 1 – Distribution of devices per industry according to Forescout Device Cloud data. Query on April 25, 2023

2. OT:ICEFALL in retrospect

2.1. Security design

While our **OT:ICEFALL** research showed the continuing prevalence of insecure-by-design practices in OT products, it also highlighted that the security controls that *did* exist were often broken – either in implementation or at the fundamental design level – while still achieving some level of security certification. This is worrying because as OT products start implementing security controls and end up getting certified, the perception of their security posture might change and the sense of urgency around compensating controls might drop – leading to a false sense of security.

Apart from typical implementation flaws such as plaintext and hardcoded credentials or key material, we encountered several recurring design issues that demonstrated many OT vendors lack a fundamental understanding of basic security control design principles – something that has also stood out in **prior work**. We detail these design issues below.

Design issue #1: Client-side authentication (CWE-603)

This is the classical example where session and authentication are completely decoupled. All authentication activity takes place in a client-side context and an attacker can simply skip this by interfacing directly with the service that is unauthenticated under the hood.

CVE-2022-33139 in Siemens WinCC OA SCADA is a clear example of this. WinCC OA components communicate through the unauthenticated, proprietary PVSS protocol, which is a legacy inheritance from the ETM PVSS SCADA system. As shown in Figure 2, WinCC OA wraps this protocol in TLS and exposes it to a web- or desktop HMI client through a proxy. Prior to version 3.17, the default authentication mode on the client was client-side authentication (CSA) where, as the name implies, credentials are fetched from the SCADA database and validated locally. An attacker can bypass authentication here by simply speaking PVSS directly to the proxy interface. While CSA is no longer the default since version 3.17, it is preferred for single sign-on (SSO) integration so continues to be deployed despite stronger server-side authentication (SSA) and Kerberos mechanisms being available.

Prior work discovered a very similar issue (CVE-2019-18250) affecting ABB Power Generation Information Manager (PGIM), an optional part of the 800xA DCS, again in a context of Active Directory (AD) integration.

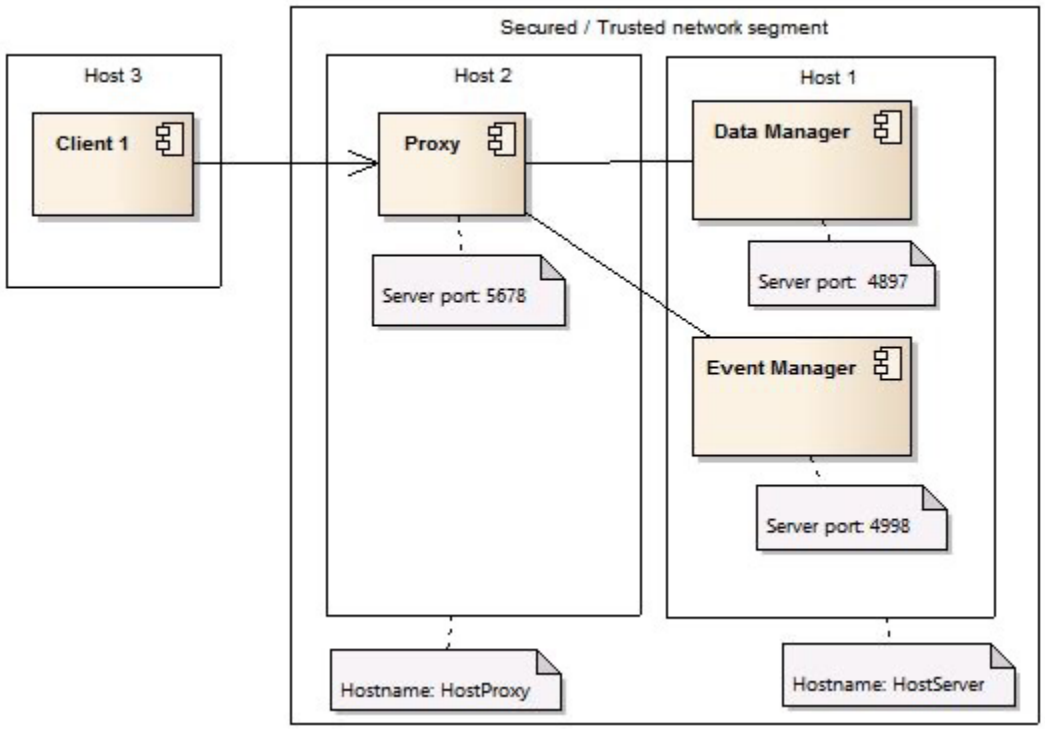

Figure 2 – WinCC OA architecture (source: WinCC OA help)

Design issue #2: Stateful controls on stateless protocols (CWE-290)

This design issue, which is exemplified by the authentication mechanisms of the Emerson ControlWave BSAP/IP protocol (CVE-2022-29954) and the Honeywell SBC S-Bus protocol (CVE-2022-30319), involves using security controls that rely on a protocol state for stateless protocols.

In both examples, successful authentication with a password will result in the client being whitelisted based on its MAC and/or IP address and all further traffic from that whitelisted entry being considered authenticated until timeout. Both protocols are UDP-based, however, and therefore stateless and allow for trivial spoofing of whitelist indicators.

Design issue #3: Missing critical step in authentication (CWE-304)

Another fundamental mistake is performing server-side authentication but then not binding subsequent traffic to that initial authentication.

A good example of this is **CVE-2022-45789** affecting **Schneider Electric Modicon *Unity* PLCs**. These PLCs use the **proprietary UMAS engineering protocol** which used to be unauthenticated and then went through several iterations of client-side authentication (**CVE-2021-22779**) until landing on the design shown in Figure 3.

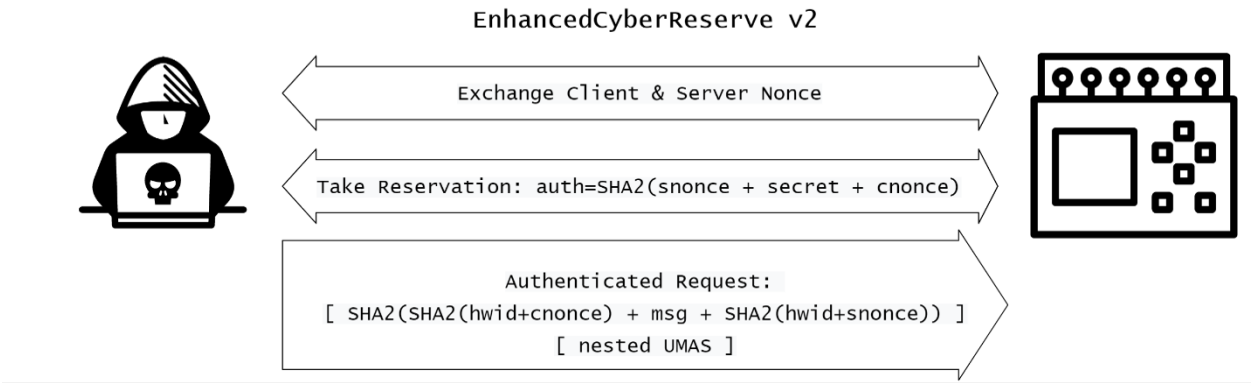

Figure 3 - Schneider Electric UMAS EnhancedCyberReserve mechanism

The *EnhancedCyberReserve* mechanism is an attempt at implementing **challenge-response authentication** that makes several core mistakes as shown in Figure 4:

1. There is no secret involved in "signing" authenticated requests; they are bound to the public nonces rather than the secret key.
2. Nonces are global and not tied to a particular session, so an attacker can sniff them, skip nonce exchange and **reuse a nonce (CWE-323)** as part of a reservation replay or request forgery attack.
3. There is no per-request freshness, allowing attackers to replay authenticated requests.

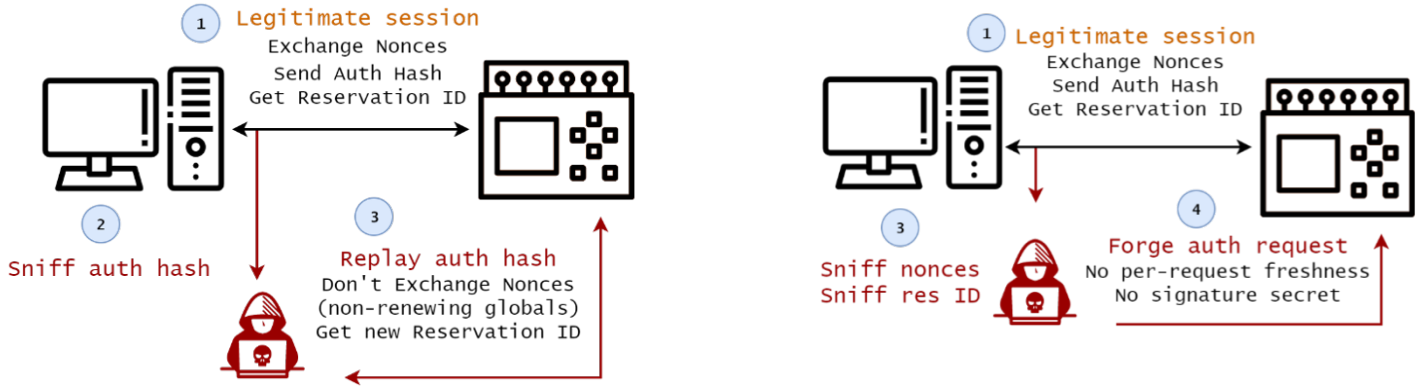

Figure 4 - CVE-2022-45789 affecting Schneider Electric UMAS

Design issue #4: Broken algorithms (CWE-327) and incorrect implementations (CWE-303)

In addition to issues with authenticated session management as discussed above, it is common for OT vendors to make basic mistakes in the authentication algorithms themselves.

This can be the use of insecure algorithms for security purposes, such as using CRC-16 for password hashing and then **using that password hash instead of a password (CWE-836)** as in Honeywell SBC S-Bus (**CVE-2022-30320**), or may involve rolling their own heavily flawed "homebrew" encryption algorithms as with the DeltaV DCS controllers (**CVE-2022-29965**).

In other cases, the issues arise from **the series of traps surrounding algorithm implementation and usage**, such as the use of the (notorious) ECB mode of operation for block ciphers in CODESYS V3 (**CVE-2022-4048**) and Motorola MDLC (**CVE-2022-30273**). Other examples in this category are the Emerson ControlWave BSAP/IP *Secure Mode* **and** *Secure Mode 2* **authentication modes** (CVE-2022-29955, CVE-2022-29956), which show a fundamental misunderstanding of challenge-response authentication: the PLC sends a key to the client, which responds with its password encrypted under that key, allowing an attacker to trivially decrypt the response.

2.2. Patch quality

The OT:ICEFALL research also exposed high variability in patch quality. While patches typically go through stringent quality assurance to ensure the system's continued functionality and stability, the *quality of the patch itself* is often less scrutinized.

In many cases, a proof-of-concept (POC) exploit delivered by researchers during disclosure is used as a unit test, which can result in the patch being tailored to a particular exploit technique rather than tackling the vulnerability's root cause. In cases where a POC is not present, either due to general absence or because tailoring a POC to a particular target is too labor intensive, the mitigation quality of a patch might not be tested at all. In extreme cases, a patch may address some vulnerabilities disclosed by researchers while **leaving others unfixed**.

Incomplete patches

This situation is prevalent in both IT and OT and can lead to *incomplete patches. Google's Project Zero defines complete patches as patches that are both correct and comprehensive:* Correct patches fix bugs with complete accuracy (no longer allowing exploitation through any route) and comprehensive patches apply fix everywhere to cover all variants.

In 2020, **Project Zero** observed that 25% of 0-days spotted that year were closely related to previously disclosed vulnerabilities, indicating the prevalence of incomplete patches, an observation echoed by **Rapid7**. In 2022, **that number** increased to 50%. **Trend Micro's Zero Day Initiative (ZDI)** made similar observations, showing a sharp decline in security patch quality in recent years, including incomplete or even placebo patches. In response, the ZDI revised their disclosure guidelines by setting a shorter 30-day deadline for vulnerabilities resulting from incomplete patches.

This is troubling because incomplete patches actually *increase* risk. Beyond leading to a false sense of security, they enable attackers in a very cost-effective way. Any patch serves as a vulnerability notification to attackers, who can **perform differential analysis** to discover the vulnerability and analyze whether anything slipped through the cracks. Well-known examples of incomplete patches leading to "new" vulnerabilities include **Log4j**, **Zoom**, **Cisco** and **SonicWall**, as well as Windows **Shell Shortcut Parser** and **Print Spooler** vulnerabilities related to flaws originally exploited by Stuxnet (respectively 5 and 10 years later!).

In our OT:ICEFALL research, we have noticed several instances of incomplete patches in OT as well, where the root cause was not properly addressed. As discussed earlier, **CVE-2022-45789** in the Schneider Electric UMAS protocol shows that **SEVD-2021-194-01** provided an *incorrect patch for* **CVE-2021-22779**. Similarly, the Emerson ControlWave BSAP/IP authentication mechanism started out with plaintext credentials and then underwent two iterations both of which (CVE-2022-29955 and CVE-2022-29956) were *incorrect patches.*

Worryingly, when such incomplete patches originate in software supply-chain components they can make their way into the products of OEMs. This occurs with Siemens **NTP** and **SQLite**, Mitsubishi **OpenSSL** and BD **Busybox** patches. In each case, incomplete patches for one vulnerability led to "new" vulnerabilities. This trickle-down effect is incredibly problematic, as it both induces a false sense of security and wastes precious security resources among vendors, asset owners and security analysts.

Interestingly, while Secure Development Lifecycle Assurance (SDLA) standards, such as **IEC 62443-4-1**, explicitly require vendors to ensure patch completeness, certified organizations continue to produce – and integrate! – incomplete patches. Whether this is incidental or systemic warrants further quantitative analysis, but these observations do cast doubt on current vendor processes.

Lack of variant analysis

One aspect of incomplete patches that deserves a special mention in OT is the general lack of **variant analysis** and the resulting *incomprehensiveness* of patches, despite SDLA requirements such as in IEC 62443-4-1, which states that "a process shall be employed for analyzing security-related issues in the product to include: [...] identifying all other products/product versions containing the security-related issue; identifying the root cause of the issue; and identifying related security issues."

In both our own research and prior work, there are many cases of OT product vulnerabilities where patches consist of piecemeal fixes.

CVE-2022-30997, concerning hardcoded credentials in Yokogawa STARDOM PLCs, is an example of this issue, considering there have been multiple prior CVEs (**CVE-2018-10592**, **CVE-2018-17896**) for hardcoded credentials in the same product on the same interface that got piecemeal patches but no variant analysis follow-up.

Other examples are **CVE-2022-29962**, **CVE-2022-29963** and **CVE-2022-29964** affecting Emerson DeltaV controllers. These pertain to hardcoded credentials that were apparently not addressed through variant analysis after similar issues (such as **CVE-2014-2350** and the silently patched **CVE-2022-29965**) had surfaced in other parts of the system.

One final, debatable example is the patch for **CVE-2022-45788** (which we discuss in our **Deep Lateral Movement in OT report**) introduced in Schneider Electric's **M340 firmware version 3.51**. Here the root cause of the vulnerability is correctly eliminated (by disabling the vulnerable command entirely) but not comprehensively addressed (since similar issues might still be present in other commands in the same subsystem that were not changed). That is not to say that this is a bad patch, merely that this kind of piecemeal patching is not the way forward in our opinion.

Slow drip patching

Next to patch completeness, an important quality factor is patch *timeliness*. Vendors and asset owners must be able to develop, ship and roll out patches as quickly as possible. Patching urgency is influenced by many factors (such as potential impact, observed in-the-wild exploitation and exposure) and must be balanced with other concerns (such as operational downtime cost and safety and stability impact analyses). Ultimately, the asset owner should be the arbiter of these decisions and should have high-quality patches available from the vendor as soon as possible, regardless of whether they decide to patch sooner or later.

Getting timely vendor responses continues to be a problem, especially when dealing with **software supply chain vulnerabilities**. In our **Project Memoria lookback report** we observed that only 22.5% of potentially affected vendors issued public advisories (as of 2022), with less than 10% of the actually affected vendors publishing patches. The average time taken by vendors affected by the vulnerabilities in Project Memoria to issue an advisory was slightly over 100 days after the public disclosure. In some cases, vendors took as long as 300 or even 588 days after public disclosure to publish advisories or patches.

The response to OT:ICEFALL disclosures improved on Project Memoria, with all vendors issuing advisories and with a better response time. OT:ICEFALL involved 61 vulnerabilities: 56 announced on **June 21, 2022,** affecting 10 vendors, three announced on **November 29, 2022,** affecting two new vendors and two announced on June 20, 2023, affecting one previous and one new vendor. Three of the original 56 vulnerabilities were not fully disclosed in 2022 at the request of one vendor, Schneider Electric. Two of these were disclosed on **February 13, 2023,** and the final one on June 20, 2023. These vulnerabilities were spread across more than 30 public security advisories issued by the different vendors. The average time to publish an advisory for OT:ICEFALL was 75 days after public disclosure and 178 days after initial notification, as shown in Table 3.

Table 3 – Time to issue security advisory for OT:ICEFALL vulnerabilities

VENDOR	DATE OF SECURITY ADVISORY	DAYS AFTER INITIAL NOTIFICATION (MARCH 10, 2022)	DAYS AFTER PUBLIC DISCLOSURE (JUNE 21, 2022)
JTEKT, Phoenix Contact and Siemens	June 21, 2022	103	0
Yokogawa	June 23, 2022	105	2
Motorola and OMRON	June 28, 2022	110	5
Bently Nevada	July 7, 2022	119	16
Emerson (DeltaV)	July 14, 2022	126	23
Honeywell (Safety Manager and Saia Burgess)	July 26, 2022	138	35
Emerson (ControlWave, OpenBSI and ROC800)	August 9, 2022	152	49
Honeywell (ControlEdge, Experion, IC protocol)	August 30, 2022	173	70
Emerson (PACSystems)	September 26, 2022	200	97
Schneider Electric (Modicon)	January 10, 2023	306	203
Schneider Electric (ION protocol)	May 9, 2023	425	322
Emerson (Ovation)	To be published	To be published	To be published
Average	**N/A**	**178**	**75**

Although better than Project Memoria on average, there were still extreme cases in OT:ICEFALL:

- Schneider Electric took 322 days to issue an advisory related to products using the ION protocol.
- Emerson has not yet published advisories for products affected by CVE-2022-29966 and CVE-2022-30267.
- Motorola did not respond to our disclosure; that vulnerability disclosure was fully handled by CISA.

Festo, CODESYS and WAGO are not included in Table 3 because they were in separate disclosure processes, which involved separate notification dates. These three vendors issued advisories on the same days as the public announcements.

Note that when a vendor issues an advisory, it does not always mean a patch is ready. For instance, Schneider Electric issued **an advisory** on January 10, 2023, for CVE-2022-45788, but the respective patch for the Modicon M340 was only made available on April 3, 83 days later. Similarly, CODESYS released **an advisory** for CVE-2022-4048 on November 23, 2022, but took 21 more days to release a patch.

Silent patching

Another phenomenon adversely affecting patch timeliness is *silent patches*, where a vendor fixes a vulnerability without publishing an advisory or registering a CVE. Sometimes silent patches are addressed with a single, non-descriptive line in a release notes file. In other cases, the issues remain unmentioned at all. As a result, OEMs and end users tend to remain ignorant of these issues and do not incorporate these patches into their products.

One example of this problem is the **bug collision** we had during our NAME:WRECK research (part of Project Memoria), where we rediscovered a 5-year old issue in Wind River IPnet (**CVE-2016-20009**) that was originally discovered by **Exodus Intelligence** but did not get a CVE. As a result, this issue lingered unaddressed for years in products such as **ABB controllers**, **BD Alaris infusion pumps**, **GE healthcare** devices, **Rockwell PLCs** and **Siemens gas turbines**. Unfortunately, other issues from this piece of Exodus Intelligence research remain without CVEs and therefore are unlikely to receive downstream patches.

In our OT:ICEFALL research, CVE-2022-29956 affecting Emerson ControlWave BSAP/IP is an example of a vulnerability resulting from a patch that is both *silent* and *incomplete*. The initial *Simple Mode* authentication was trivially broken and led to the introduction of the publicly documented *Secure Mode*. While this mechanism was trivially broken as well, it did not receive a CVE until we reported the issue (CVE-2022-29955) and was patched silently with the introduction of the undocumented *Secure Mode 2*. Because of this, known issues in (at least) the regular *Secure Mode* remained under the radar for years, while issues in *Secure Mode 2* were not discovered due to lack of public advisory.

2.3. Vendor security testing

SDLA standards such as IEC 62443-4-1 require vendors to perform vulnerability testing through “manual or automated abuse case testing,” explicitly mentioning fuzz testing as one example. Similarly, product security certification programs such as those by **GE Achilles** explicitly claim grammar-based protocol stack testing is part of the program (as shown in Figure 5). Terminology aside, clearly the intention behind both is to conduct at least basic fuzz testing of OT product protocol stacks.

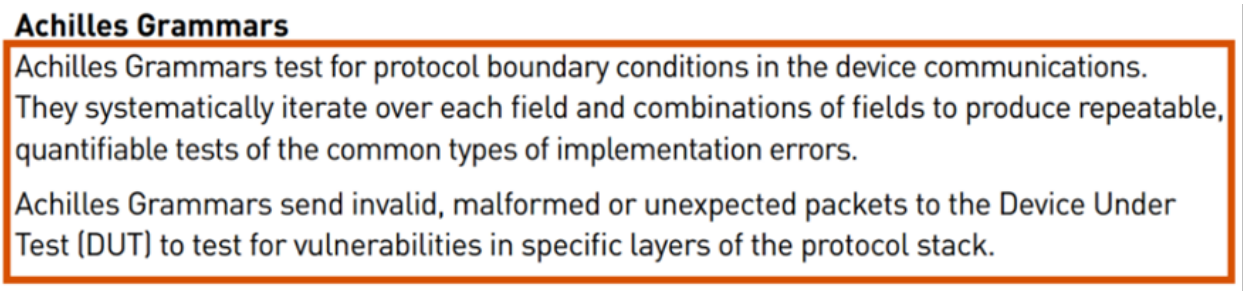

Figure 5 – Vulnerability testing requirements in GE Achilles Certification

It was somewhat surprising to find that many products affected by vulnerabilities such as those in Project Memoria, Urgent/11 and Ripple20 were both certified at a product level and developed by organizations with certified secure software development life cycles (SDLCs). For example, Mentor Graphics (now Siemens) Nucleus RTOS, Wind River VxWorks 7 RTOS, Siemens SENTRON PAC4200 power meters, Schneider Electric ATV6000 variable speed drives, ABB AC 800M PM865 DCS controller and Rockwell ControlLogix 1756-EN2* communications modules are all GE Achilles Communications Certified (ACC) Level 2 products. Considering the shallow nature of the vulnerabilities involved, this casts serious doubt on the quality of the fuzz testing these products are supposed to have undergone.). Terminology aside, clearly the intention behind both is to conduct at least basic fuzz testing of OT product protocol stacks.

It was somewhat surprising to find that many products affected by vulnerabilities such as those in **Project Memoria**, **Urgent/11** and **Ripple20** were both certified at a product level and developed by organizations with certified secure software development life cycles (SDLCs). For example, **Mentor Graphics (now Siemens)** **Nucleus RTOS**, **Wind River VxWorks 7 RTOS**, **Siemens SENTRON PAC4200 power meters**, **Schneider Electric ATV6000 variable speed drives**, **ABB AC 800M PM865 DCS controller** and **Rockwell ControlLogix 1756-EN2*** **communications modules** are all GE Achilles Communications Certified (ACC) Level 2 products. Considering the shallow nature of the vulnerabilities involved, this casts serious doubt on the quality of the fuzz testing these products are supposed to have undergone.

In addition, at least two vendors in those examples had IEC 62443-4-1 SDLA certification. The products in question were developed before these organizations achieved SDLA, which leads one to believe that the certified SDLC best practices are only being applied to products manufactured after certification.

3. Conclusion

The findings in OT:ICEFALL demonstrate the need for tighter scrutiny of and improvements to processes related to security design, patching and testing in OT device vendors.

3.1. Secure SDLC maturity

OT security has reached a state where there are increasing international discussions about the need for more **vendor liability** and better **security by design and by default**. One of the strategic objectives in the US 2023 **National Cybersecurity Strategy** is to "shift liability for insecure software products and services," which would entail legislation to establish liability of device vendors for insecure or vulnerable products. Similarly, the EU is working on a **Cyber Resilience Act**, which has as its first goal to "ensure that manufacturers improve the security of products with digital elements since the design and

development phase and throughout the whole life cycle." As a final example, the UK's recently passed **Product Security and Telecommunications Infrastructure Act** has specific sections about IoT security requirements, duties of manufacturers and potential penalties.

Regardless of how these regulatory discussions will evolve or how the laws will be enforced, a good way to start improving the overall state of OT security would be to ensure that vendors address obvious design flaws such as the ones outlined in Section of this report – and others that we did not dive into, such as the continued use of hardcoded credentials – both before a product goes to market and as part of continued support for existing products.

Shifting security efforts left (to guarantee application security at the earliest stages in the development cycle) will also break the current culture of inefficient and disruptive piecemeal patching in OT. Patches aren't free or without risk in OT; the patch cycle ROI should be maximized as much as possible. Theoretically, this would be addressed through certified secure SDLCs and product security certifications. As we have pointed out, however, in practice these might not be mature enough yet. Periodic external scrutiny (both commissioned by vendors themselves and by asset owners during procurement) will likely remain necessary to raise the bar.

An important second step is to address incomplete and late patches, as we detailed in Section . Incomplete patches have been identified as a **source of vulnerabilities** for decades and can affect IT, OT and IoT products. There are currently **close to 1,000 vulnerabilities** specifically mentioning they exist because of an "incomplete fix," as well as countless others that exist for the same unnamed reason. **Recent research** shows that faulty patches cause delays in improving security in organizations as critical as hospitals. Ensuring that patches undergo strict security testing with variant analysis and are given priority over new product features would automatically decrease the number of new vulnerabilities. Again, this quality assurance should be a given by certified secure SDLCs but in practice too often is not.

3.2 Vulnerability management

The observations above have implications for vulnerability management in OT. Even with secure designs, secure default configurations and improved patch quality assurance, serious vulnerabilities are bound to be discovered from time to time and patches will need to be applied. However, **patching in OT is unlike patching in IT**, where patches tend to be rolled out indiscriminately across assets during regular patching cycles. In OT, patching is a complicated, time-consuming and potentially risky process. First, vendor patches must be approved and tested by the OEM and system integrator, after which asset owners must walk through the entire patch path, with backlogs sometimes spanning years. Before applying a patch, backups of old software and firmware versions, configurations and programmable logic must be made for rollback in case of issues, and dependencies must be mapped across assets – patches may contain a mix of security and functional updates requiring other assets to also be upgraded to maintain compatibility. On top of that, the most critical assets are sometimes caught in the **Cyber Maintenance Catch-22**.

Persistence of insecure-by-design in OT and subpar patch quality assurance tend to mean that many security patches deliver minimal risk reduction at significant cost. As a result, some asset owners contend that patching OT vulnerabilities is futile and defenders should focus on compensating controls such as segmentation and network- and host monitoring. While we agree asset owners should carefully consider when and how to patch in OT (if at all), there are some nuances to be observed.

Consider an IED in an electrical substation for which circuit breakers can be opened and closed through the insecure-by-design **IEC 61850 MMS** protocol, which is not carried over TLS tunnels or restricted by IP access control list. An authentication bypass on its webserver that allows attackers to manipulate breakers is almost irrelevant in this case. However, a DoS vulnerability affecting the protective functionality and requiring a power cycle to reset (e.g., **CVE-2015-5374**) or, even worse, an unauthenticated firmware download that allows an attacker to brick the IED, are very different. In the former case, defenders can close an opened breaker remotely via the HMI or engineering tools after evicting the attacker. In the latter cases, breakers will have to be closed manually and IEDs will need to be power-cycled or even replaced, resulting in much longer outages and harder to scale incident response.

In cases like this, where a vulnerability allows attackers to achieve a different consequence than possible through normal functionality, patches can reduce risk more efficiently than (non-cyber) engineering controls would.

People who favor a **consequence-driven approach** will likely prioritize such issues, whereas those more concerned about the likelihood of an attack will 1) remove assets from the public-facing internet and 2) patch their domain controllers and multi-homed servers still running legacy Windows – because that is what is mostly targeted by opportunists and ransomware gangs.

Even so, organizations with more mature OT security will eventually reach a point where such low-hanging fruit will only exist deep down in the networks of specific remote sites. Once those assets are secured, patching priorities need to increasingly consider the mission-criticality of an asset or its connectivity to mission-critical assets.

Consider, too, that attack likelihood is not a constant. It is influenced by factors such as attacker motive and publicly available capabilities. For example, **EternalBlue** was the first public remote Windows exploit in years when it leaked and was only accessible to top-tier attackers before rapidly becoming available to anyone who could run Metasploit. Similarly, the **CVE-2015-5374** Siemens SIPROTEC DoS was first reported (without details) in **July 2015** before being used in December 2016 as part of the **Industroyer** attack. Perhaps the lack of details made the exploit less likely to be used, but they can be found in a **presentation** from the May 2016 PHDays VI conference in Moscow, where a 17-year old highschooler used it in a CTF against a model electrical substation – *six months* before the Industroyer attack. In March 2018, the exploit **was integrated** into the Metasploit framework, rendering it available to the widest possible audience. Similar OT modules in Metasploit for other protocols and devices have also been used by **opportunistic attackers**.

The principle of defense-in-depth is designed precisely to deal with this kind of "likelihood volatility," which is especially important given the inflexibility of OT patching. If a defender banks on low likelihood alone, they might not be able to patch quickly enough if that likelihood suddenly changes (e.g., when yet another **firewall or VPN vulnerability** breaks those segmentation-only defenses).

Based on the abundant evidence gathered and insights gained throughout OT:ICEFALL and summed up in this report, we consider a careful, consequence-driven analysis of which vulnerabilities to patch, in which assets, to be the best approach for organizations with more mature OT security programs. By prioritizing issues based on potential consequences, you will reduce security risk and minimize operational downtime more effectively than by relying too heavily on vendor guidance or compensating controls.

www.forescout.com/research-labs/ **vederelabs@forescout.com**

FORESCOUT

Forescout Technologies, Inc.
Toll-Free (US) 1-866-377-8771
Tel (Intl) +1-408-213-3191
Support +1-708-237-6591
Learn more at Forescout.com