FORESCOUT

Concluding Project Memoria – Lessons Learned after 18 Months of Vulnerability Research

By Forescout Research Labs

Table of Contents

Disclaimer: In this report, we account for both vulnerabilities clearly branded as Project Memoria (AMNESIA:33, NUMBER:JACK, NAME:WRECK, INFRA:HALT and NUCLEUS:13), as well as vulnerabilities strictly connected to Project Memoria, such as Ripple20.

1. A Summary of Project Memoria

The idea of Project Memoria emerged in May 2020 while collaborating with JSOF on Ripple20. Our researchers immediately understood that the problem with TCP/IP stacks was much deeper and much more widespread than initial research had suggested. We hypothesized that similar issues to those identified in Ripple20 could be present in other stacks as well – many other stacks, distributed in many flavors across many vendors and many products.

The will to validate this hypothesis led to our first study, AMNESIA:33, where we confirmed that many open-source TCP/IP stacks shared several similar vulnerabilities. Building on top of the knowledge gathered with AMNESIA:33, we decided to extend the project by first looking into the presence of similar bugs in different stacks (NUMBER:JACK and NAME:WRECK) and then looking into security flaws of very specific stacks (INFRA:HALT and NUCLEUS:13).

There are a total of **97 vulnerabilities under Project Memoria**: 19 in Ripple20 found by JSOF, 33 in AMNESIA:33, nine in NUMBER:JACK, nine in NAME:WRECK found together with JSOF, 14 in INFRA:HALT found together with JFrog, and 13 in NUCLEUS:13 found together with Medigate. **These vulnerabilities affect 14 TCP/IP stacks**: CycloneTCP, FNET, FreeBSD, IPnet, MPLAB Net, NetX, NicheStack, NDKTCPIP, Nucleus NET, Nut/Net, picoTCP, Treck, uC/TCP-IP and uIP. lwIP remains the only stack we analyzed and did not find any issue.

Figure 1 summarizes some of the main achievements of the project. Below, we will discuss the main lessons we learned from Project Memoria.

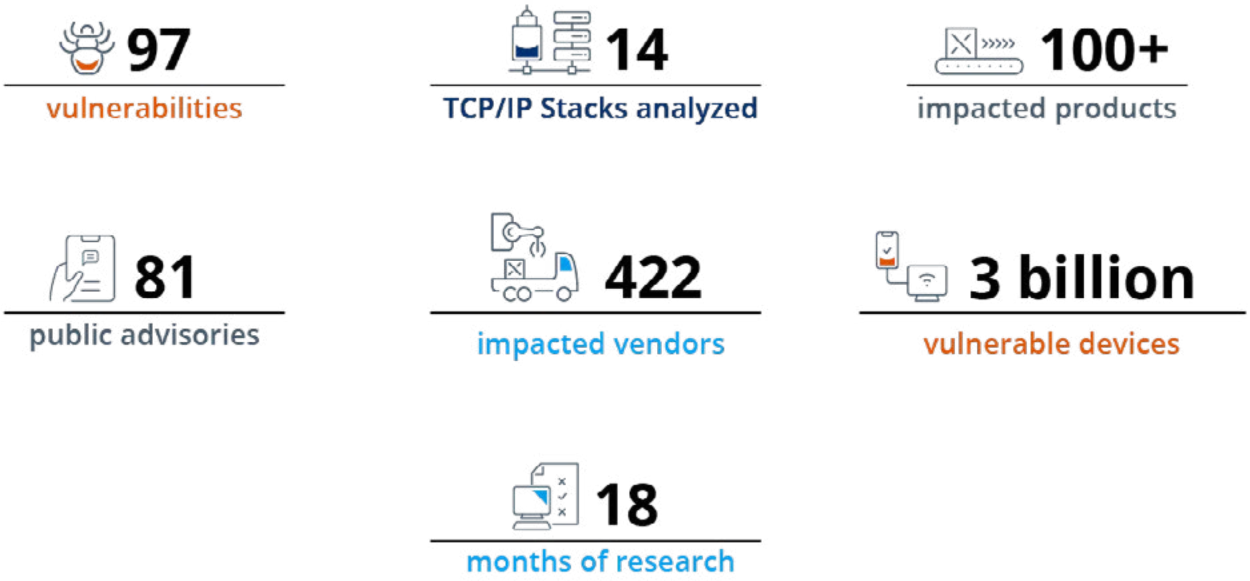

Figure 1 – Project Memoria Summary

2. The Lessons Learned

2.1. Legacy Software Enables the Connected World

TCP/IP stacks have been around for a long time, and they have a variety of decades-old vulnerabilities, which often affect different versions of a stack. Table 1 shows the year (and age) of initial release of each of the stacks investigated under Project Memoria. The newest stack is seven years old, while the oldest are 28. The average age is 18.85 years – almost two decades. Clearly, these stacks were originally designed and implemented at a time when cybersecurity was not as big of a concern as it is today.

Although these stacks are still actively developed, it is common that some vulnerabilities that have been patched by the stack vendor do not make it all the way down the supply chain to all the affected devices. Below we discuss one of the main reasons: silent patching.

Stack Name	Year of Initial Release	Age of Initial Release
CycloneTCP	2013	8
FNET	2003	18
FreeBSD	1993	28
IPnet	? (acquired by Wind River in 2006)	15
MPLAB Net	2014	7
NetX	1997	24
NicheStack	1996	25
NDKTCPIP	2010	11
Nucleus	1993	28
Nut/Net	2002	19
picoTCP	2013	18
Treck	1997	24
uC/TCP-IP	2002	19
uIP	2001	20

Table 1 – Year of initial release of TCP/IP stacks

2.2. Silent Patching Is a Terrible Idea

Silent patching refers to the practice of fixing a vulnerability without public documentation and without assigning a CVE ID. This has always been common practice among software vendors and is slowly changing, with some vendors becoming more open to assigning CVE IDs to issues that are internally discovered or that affect older versions of their software.

In Project Memoria, we first encountered silent patching with NAME:WRECK. CVE-2016-20009 was originally discovered by Exodus Intelligence in 2016 and never assigned a CVE ID; we independently found the issue again in 2020 and spent months (with the help of the CERT/CC) convincing Wind River – the owners of IPnet/VxWorks – to assign an ID to the issue. After the publication of NAME:WRECK, as downstream vendors became aware of this five-year-old patched issue, there have been several vendor advisories listing critical vulnerable devices, such as ABB controllers, BD Alaris infusion pumps, GE healthcare devices, Rockwell PLCs and Siemens gas turbines.

In NUCLEUS:13, we again saw instances of silently patched vulnerabilities, although this time, Siemens was very proactive in assigning the CVE IDs. We do expect that other downstream vendors will again re-evaluate their products and see that there are vulnerabilities they were not aware of.

Project Memoria shows two things about silently patched vulnerabilities:

1. They exist in very critical supply-chain software, so there are millions of devices out there that have been vulnerable for a long time without even their vendors knowing about it because other vendors chose to remain silent.

2. Silently patching a vulnerability does not mean that nobody will get to know about it: these issues tend to be rediscovered again and again.

2.3. Vulnerabilities Are Predictable (and We'll Tell You Why)

A clear finding of Project Memoria is that bugs are almost predictable. We have seen the same mistake happening repeatedly. We have shared with the community the knowledge we gathered about '*what developers should not do*' as a set of 11 anti-patterns, summarized in Table 2.

#	Anti-Pattern	Study
1	Absence of bounds checks	AMNESIA:33
2	Misinterpretation of RFCs	AMNESIA:33
3	Shotgun parsing	AMNESIA:33
4	IPv6 extension headers/options	AMNESIA:33
5	Predictable ISN generation	NUMBER:JACK
6	Lack of TXID validation, insufficiently random TXID and source UDP port	NAME:WRECK
7	Lack of domain name character validation	NAME:WRECK
8	Lack of label and name lengths validation	NAME:WRECK
9	Lack of NULL-termination validation	NAME:WRECK
10	Lack of record count fields validation	NAME:WRECK
11	Lack of domain name compression pointer and offset validation	NAME:WRECK

Table 2 – Identified anti-patterns

These anti-patterns appeared in other vulnerabilities found as a direct or indirect result of the work in Project Memoria, such as CVE-2020-17528 and CVE-2020-17529 on NuttX (which are essentially the same as two uIP CVEs in AMNESIA:33 but tracked under a different number), CVE-2021-25663 and CVE-2021-25664 on Nucleus NET (found by Siemens and similar to IPv6 issues found in AMNESIA:33), as well as additional 4 CVEs on Treck found by Intel as part of their internal investigation.

By now, we are confident that continuing the project with other stacks would yield many other CVEs. At the same time, we are also confident that we have identified most of the "low-hanging fruit" anti-patterns, and that we should start doing the same for other important supply-chain components. Therefore, as we move to other targets for our research, we invite the security community to continue the work on TCP/IP stacks by using (and improving) the artifacts we have shared, which include: a fingerprinting script, a static analysis tool, a draft RFC, exploit PoCs (that can be requested by sending an email to research@forescout.com and that, so far, have been mostly privately shared with CERTs and vendors) and traffic samples (which **were shared with 15 cybersecurity vendors**, including direct competitors).

2.4. Vendors Are Often Unresponsive

Identifying vulnerable vendors and devices has been the greatest challenge under Project Memoria. In the past year, we have helped many device vendors understand the impact of the vulnerabilities, we have reviewed patches, and we have helped asset owners identify and mitigate the risks around vulnerable yet un-patched devices.

Although we expected that identifying vulnerable vendors, products and models would be a challenge, we were surprised to realize how difficult it is to keep track of vendors' responses. Often, we were surprised to find out that a vendor issued a security advisory months after our public disclosure, which we only found out about because we were proactively searching.

Although there are initiatives to coordinate and centralize vulnerability disclosure (such as the CERT/CC Vulnerability Information and Coordination Environment), there is no central communication when a new product is found to be vulnerable to Project Memoria nor a central location aggregating the data. **Most vendors often issue advisories on their websites or directly to their customers**, leaving the rest of the community in the dark. For instance, **BD has issued an advisory on their website regarding NAME:WRECK affecting their Alaris PC Unit infusion pumps**, but there is no link to our research (so that interested readers could better understand the issues), and they call the vulnerability "WRECK," which adds confusion to the matter.

To **try and mitigate this issue, Forescout Research Labs is maintaining a page with the full list of advisories connected to Project Memoria**. The page can be found here, and we invite interested parties to contribute to it.

The figure below shows a summary of vendor response throughout the project.

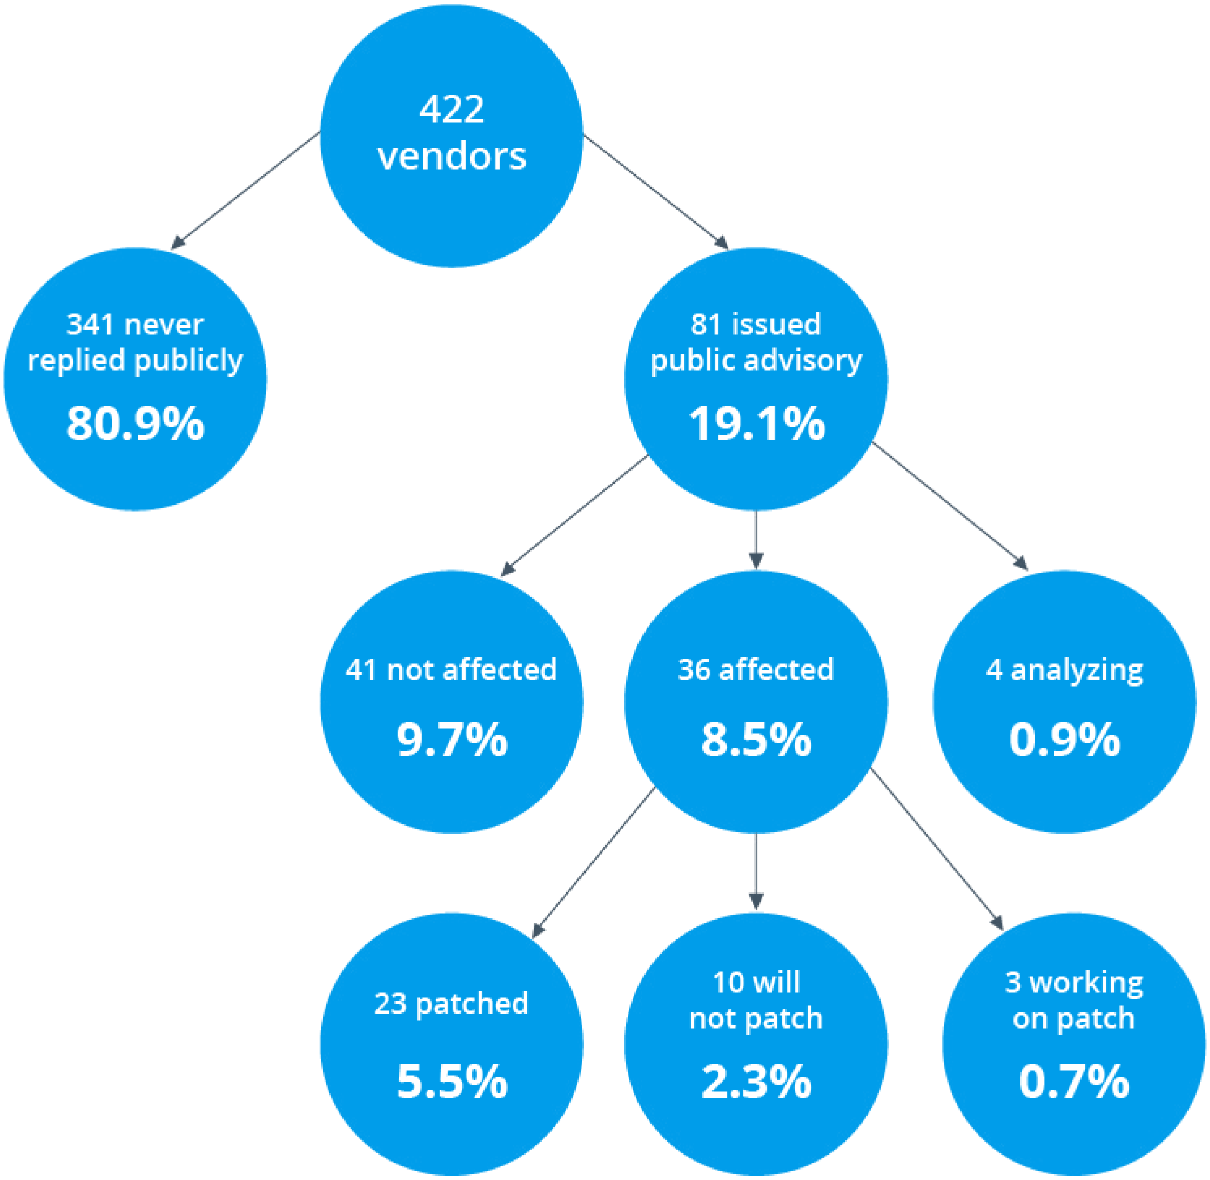

In the different phases of the project, we identified (and shared with several agencies) a total of 422 device vendors that could be using the vulnerable stacks. We currently track 81 vendors that have issued public advisories: 36 have confirmed to be affected, 41 have confirmed not to be affected, and four are still investigating.

Of the 36 affected, 10 will not provide patches, and three are still working on them. So far, **the vendor with the longest time to publish** patches was Schneider Electric: their patches for AMNESIA:33 were published on 12/October/2021, **308 days after the public disclosure** on 08/December/2020.

This means that only 19% of potentially affected vendors have provided some public response, and only 5.5% have actually patched the vulnerabilities so far. Even considering that other vendors have published private advisories (either directly to customers or to coordinating agencies that have not shared the information back to us), this highlights a huge gap in the current vulnerability disclosure and public management processes. In the end, the customers and asset owners of most of the potentially affected vendors are left with no clue on whether their devices are vulnerable or not.

Among all the vendors, Siemens is the only one that has publicly stated to be affected by the vulnerabilities in all the disclosure phases. So far, Siemens has issued 12 advisories based on Project Memoria's findings. Siemens is also the vendor that issued 31% of ICS-CERT alerts in 2020. This is not a coincidence and is far from implying that Siemens' devices are less secure than others. On the contrary, it shows that they have a mature product security program and that they are open to acknowledging and publishing issues that affect their products. It also indicates that several other similar vendors have not taken the same proactive approach and may be leaving their customers vulnerable.

2.5. Hundreds of Products Are Impacted

The TCP/IP stacks analyzed under Project Memoria are used in a wide variety of connected devices, so the potential impact of the same vulnerability spans several industries, such as healthcare, government, financial, manufacturing and transportation. Below we list some examples of products impacted (as confirmed by published advisories):

- B&R Automation motion control and track technology used in manufacturing
- Siemens gas turbines, high power transmission devices and RTUs, as well as EMU smart meters used by electrical utilities
- Ricoh printers and Extreme Networks switches used in corporate organizations
- FEIG and Siemens RFID readers used in logistics and retail
- BD infusion pumps and Philips patient information systems used in healthcare
- Microchip WiFi modules used in consumer electronics and Philips consumer products such as robot vacuum cleaners and air purifiers
- Industrial controllers from major manufacturers such as Phoenix Contact, Rockwell Automation, Schneider Electric and others used for several different functions, such as building automation, manufacturing control and electric vehicle charging stations

However, since public information about vulnerable devices coming from the vendors is very rare, that is only a subset of devices actually impacted. To understand the real impact of these vulnerabilities in organizations, we must make use of a different data source: the Forescout Device Cloud.

We identified a quarter of a million devices affected by any of the issues in Project Memoria. Figure 2 shows a breakdown of the number of vulnerable devices in five verticals: government, healthcare, manufacturing, retail and financial services. **Government and healthcare have the highest number of vulnerable devices, followed by manufacturing and retail**. Our data also shows that 67% of the organizations we track are affected by these vulnerabilities.

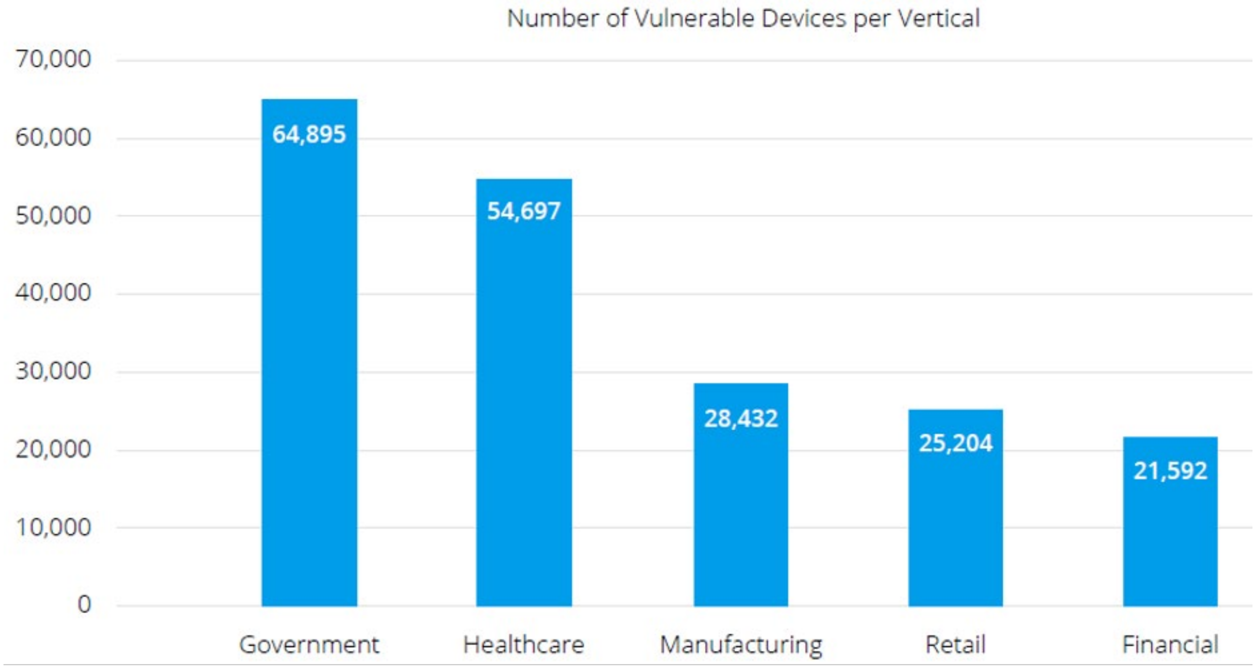

Figure 2 – Vulnerable devices per vertical

2.6. Organizations across Several Industries Are Impacted

We continue the analysis on the Device Cloud to show the impact of these vulnerabilities on organizations in different verticals.

Figure 3 shows that, on average, every organization has 200 vulnerable devices, while healthcare has by far the largest average number of vulnerable devices – almost 500 – per organization. This is partly explained by the diversity of specialized devices that are common in healthcare environments. Hospitals in the United States, for instance, had, on average, between 10 and 15 connected medical devices per bed in 2016.

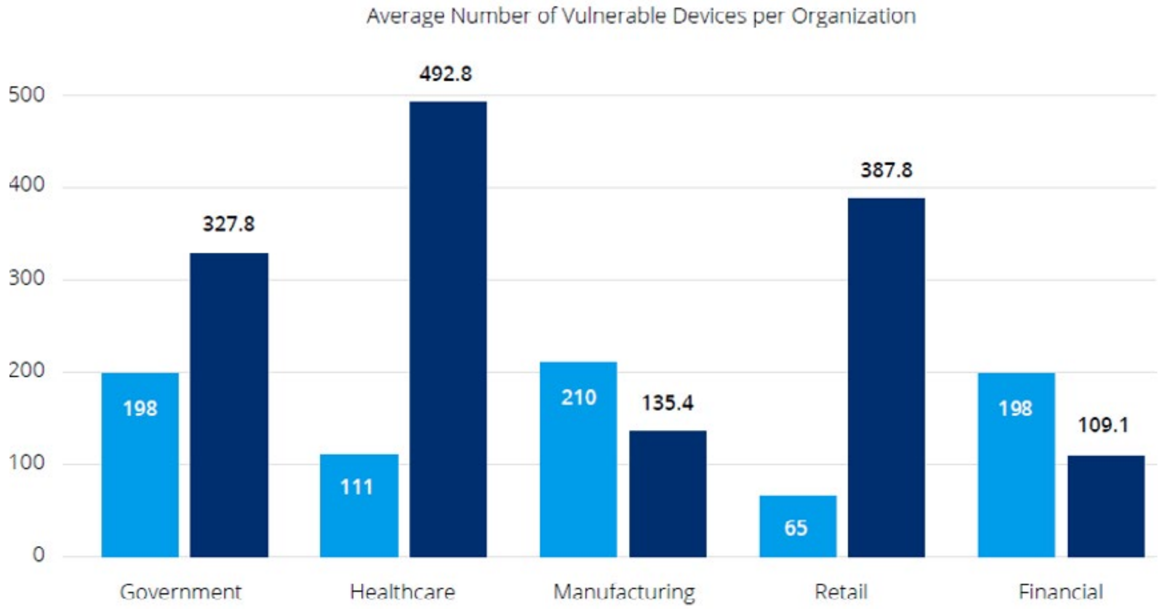

Figure 3 – Vulnerable devices per organization

Figure 4 shows a breakdown of distinct vulnerable device types per vertical, something that we call device diversity. The implication of high device diversity within an organization is that patching vulnerabilities will be more time-consuming. In networks with high device diversity, security operators must spend a considerable amount of time to identify and patch vulnerable devices. This is because (1) the tools able to identify IT devices might differ from those able to identify medical or IoT devices, and (2) with different device types come different vendors and, hence, patches available with different timelines and applicable with different procedures.

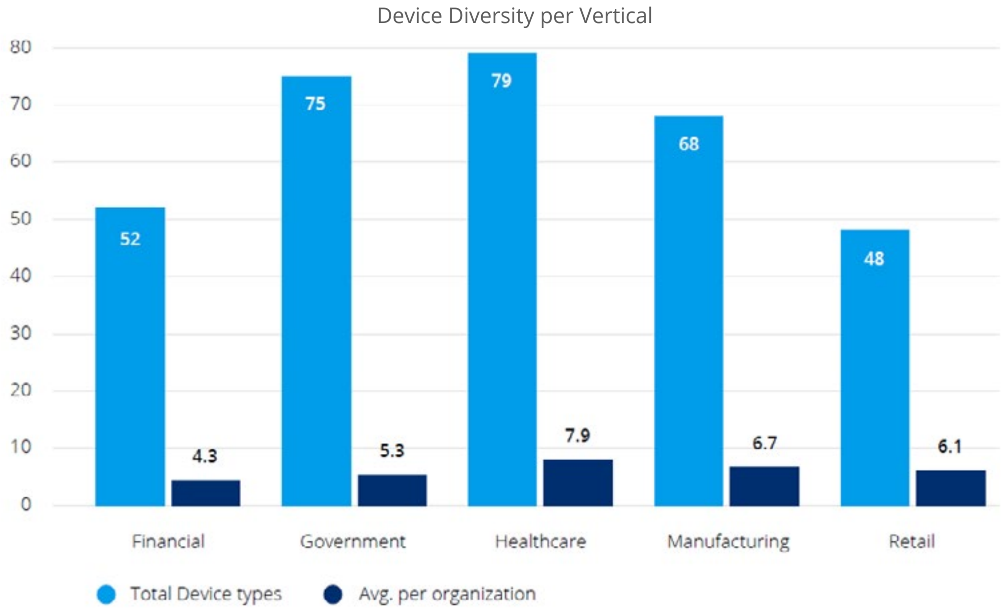

Figure 4 – Device diversity per vertical

Figure 5 shows the average number of distinct vendors affected by TCP/IP vulnerabilities, something we call vendor diversity. As for device diversity, a high vendor diversity is directly connected to more time needed to apply patches. According to the data in Figure 5 (dark blue bars), healthcare has the highest average diversity per organization (12), followed by manufacturing and retail (about 10). By looking at each vertical as a whole (light blue bars), manufacturing has the absolute highest number of vendor diversity (293 vulnerable vendors over 210 organizations), followed by healthcare (259 vulnerable vendors over 111 organizations). Since patches for TCP/IP stack vulnerabilities must trickle down the supply chain, several of those vendors either do not issue patches or take months to do so, which means the affected devices remain vulnerable for a long period of time.

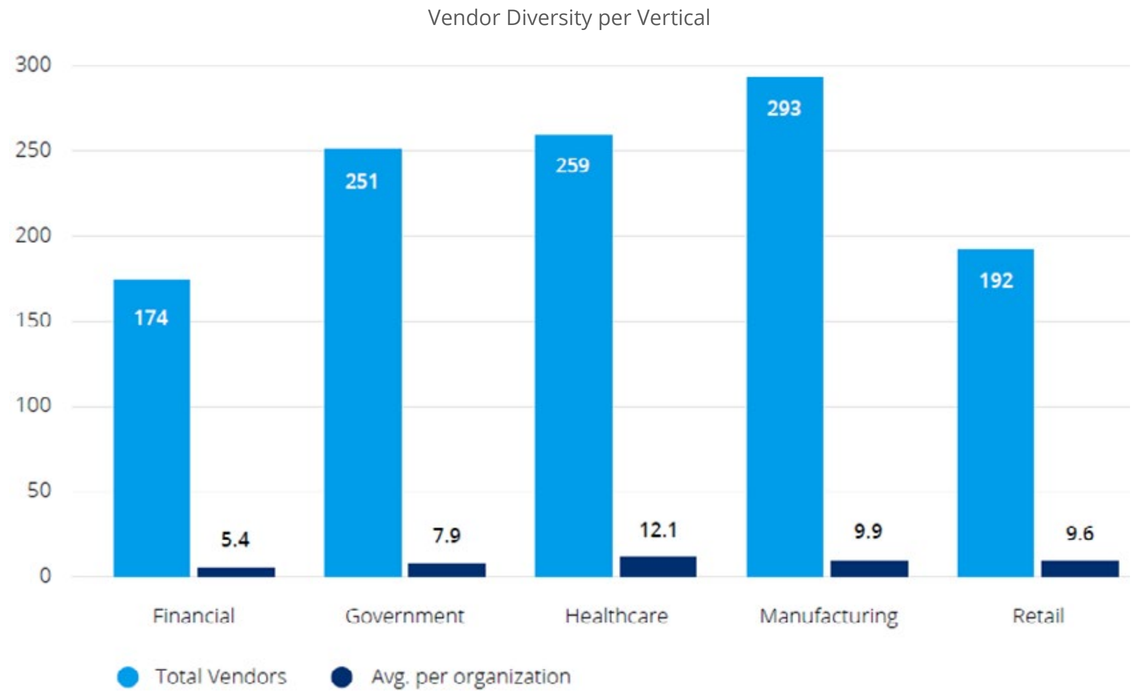

Figure 5 – Vendor diversity per vertical

Figure 6 shows that printers (34%), IP phones (20%), networking devices (8%), building automation (8%) and infusion pumps (4%) are the most common device types vulnerable to TCP/IP stack vulnerabilities in all the organizations we track. The Figure also shows the top five most common vulnerable device types in each vertical.

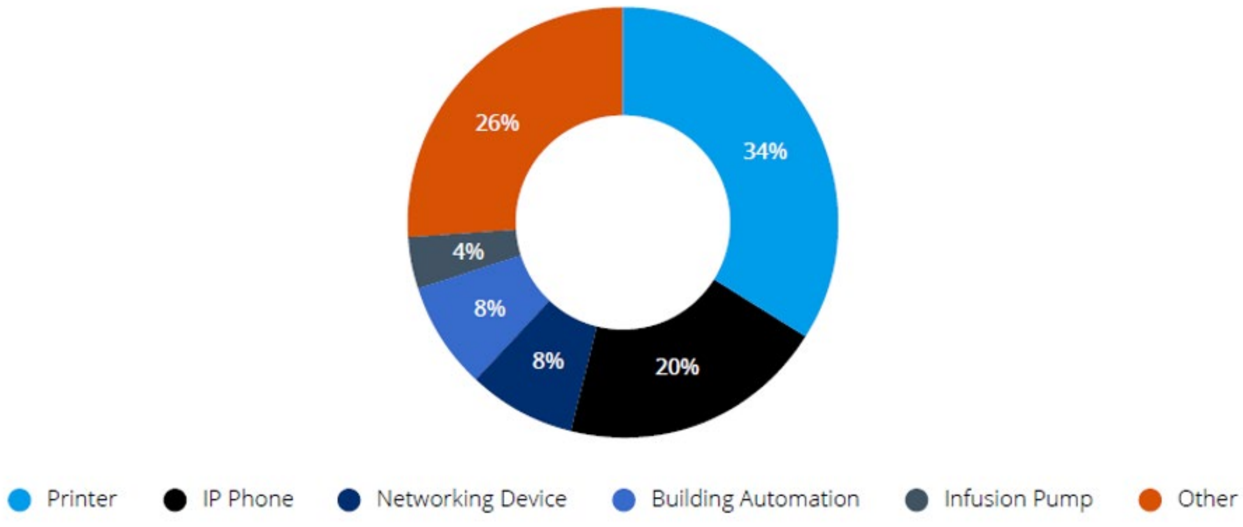

Figure 6 – Most common vulnerable device types across verticals

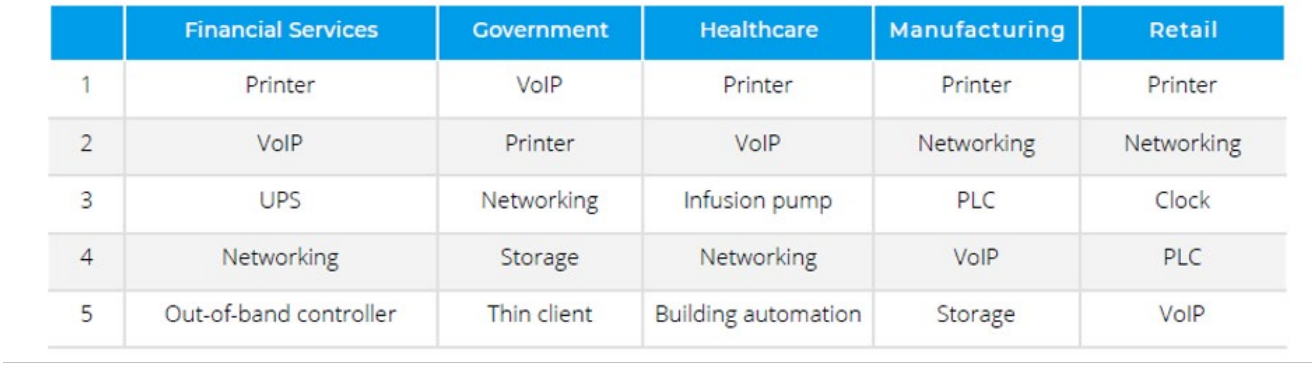

	Financial Services	Government	Healthcare	Manufacturing	Retail
1	Printer	VoIP	Printer	Printer	Printer
2	VoIP	Printer	VoIP	Networking	Networking
3	UPS	Networking	Infusion pump	PLC	Clock
4	Networking	Storage	Networking	VoIP	PLC
5	Out-of-band controller	Thin client	Building automation	Storage	VoIP

Table 3 – Top 5 vulnerable device types in each vertical

The business cost of vulnerable devices translates into the increased risk of cyberattacks. Data breaches in 2021 cost an average of more than US$4 million to organizations, which includes items such as forensic and incident response activities, legal expenses and regulatory fines, as well as the biggest cost: lost business. Vulnerabilities on TCP/IP stacks often affect OT and IoT devices that are directly connected to business operations – such as the infusion pumps and PLCs shown in Figure 6 – which means that exploiting them can cause system downtime that immediately leads to lost business. The cost of system downtime has been measured by Gartner in 2014 as $5,600 per minute across every type of organization.

2.7. SBOMs Help Mitigate the Problem

Project Memoria is about supply chain vulnerabilities and thus inherently connected to another interesting initiative, the Software Bill of Materials (SBOM).

When we published AMNESIA:33 in December 2020, we discussed the lack of Software Bill of Materials (SBOMs) and opaque supply chains as the biggest challenges to identifying vulnerable devices. While our project progressed, other researchers have done very important work on vulnerabilities affecting other important supply-chain components, such as DNS forwarders (DNSpooq) and RTOSes (BadAlloc).

At the same time, attackers have realized that compromising supply chains is an extremely effective way of targeting organizations. The past year has seen devastating attacks leveraging, for instance, SolarWinds and Kaseya to infiltrate hundreds of organizations. ENISA has recently published an extensive analysis of the threat landscape for supply chain attacks, and researchers are now analyzing several system administration tools that could be leveraged in similar attacks.

The cybersecurity community has reached a point where both industry and government recognize the complexity of software supply chains and the importance of SBOMs to fix the supply chain vulnerability problem. Luckily, there is important progress being made on that. We are proud to have been a small part of what made progress possible by highlighting and bringing awareness to an important topic.

3. Conclusion

We discussed the main lessons learned from a project that identified almost 100 new vulnerabilities across 14 TCP/IP stacks in the past year and a half. Concluding Project Memoria does not mean that our work is done, either for TCP/IP stacks or other foundational components of the connected device ecosystem. As we did in previous studies, we invite other researchers and device vendors to continue this work and collaborate with us in future research.