FORESCOUT RESEARCH | VEDERE LABS

Trends in cyberattacks, exploits, and malware

2023 Global Threat Roundup Report

January 24, 2024

Contents

Main takeaways from 2023

China escalates cyberattacks. Threat Actors across the globe are targeting web applications and remote access. Nearly half of attacks come from ISPs. IoT and network infrastructure exploits rise. CISA covers only 35% of known exploited vulnerabilities. Remote access trojans (RATs), infostealers and botnets are the most used malware types. Top threat actors are based in China, Russia and Iran. Main targets include USA, UK, Germany. Devices are exploited both in targeted and opportunistic attacks.

1. Executive summary

Our inaugural threat roundup report last year started by observing that "*the year 2022 was eventful for cybersecurity*." As you can imagine, 2023 was no less eventful. Some of the key events included ongoing geopolitical conflicts and the appearance of new ones, the emergence of critical vulnerabilities being mass exploited and the ever-increasing threat of cybercrime.

In this report, we look back at all the data we have available about attacks and the threat landscape of 2023 to share with organizations tactical insights and strategic recommendations for improved defense.

Key findings of this report include:

- **Attacks originated from 212 countries.** The top 10 countries accounted for 77% of the malicious traffic, with a spike in attacks originating from China. 48% of attacks came from IPs managed by ISPs, 32% from organizations in business, government and other sectors, and 10% from hosting or cloud providers. This reflects an increase in the use of compromised devices to launch attacks, whether directly or via "residential proxies."
- **Web applications were the most attacked service type followed by remote management protocols.** Remote management services were often targeted with specific usernames linked to IoT devices, whereas web applications were often targeted with vulnerability exploits.
- **Exploits against software libraries decreased partly because of Log4j exploits losing popularity.** Exploits against network infrastructure and IoT devices increased. The most targeted IoT devices were IP cameras, building automation and network attached storage. Only 35% of exploited vulnerabilities appeared in CISA KEV.
- **Five OT protocols were constantly targeted:** Modbus (a third of attacks), Ethernet/IP, Step7, DNP3 (with around 18% each) and IEC10X with 10% of attacks. The remaining 2% represent many other protocols, of which the majority is BACnet. Most attacks target protocols used in industrial automation and the power sector. Building automation protocols are less often scanned, but exploits against building automation are more common.
- **Post-exploitation actions focused on persistence** (50%, up from 3% in 2022), discovery and execution. Most observed commands are for generic Linux systems, but there were also commands executed specifically for network operating systems that run on popular routers.
- **We observed an equal amount of remote access trojans (RATs) and information stealers (infostealers) as the most popular type of malware.** Botnets and other downloaders come in third and fourth, followed by crypto miners and then a variety of other malware, such as keyloggers and adware. The most popular malware families observed were the Agent Tesla RAT (16%), then variants of the Mirai botnet (15%) and the Redline infostealer (10%).
- **Cobalt Strike remained the most popular command and control (C2) server** (46%), followed by Metasploit (16%) and the emerging Sliver C2 (13%). Most C2s are in the United States (40%), followed by China (10%) and Russia (8%).
- **Threat actors targeted 163 countries.** The United States was the most targeted by far, with 168 actors aiming at the country. In second place came the United Kingdom with 88, then Germany with 77, India with 72 and Japan with 66. Most threat actors were in China (155), Russia (88) and Iran (45). Together, these

three countries accounted for almost half of threat actor groups in our database. Government, Financial Services and Media and Entertainment were the industries most targeted by these actors.

- **Most attacks we observed are opportunistic.** Although, there were exploits targeting very specific networking devices to obtain precise information about them and drop malware. These attacks often use public proof-of-concept scripts.

Where does our data come from?

Most data used for the analysis in this report comes from the Vedere Labs **Adversary Engagement Environment** (AEE), a set of honeypots on the open internet luring attackers and recording their actions. The data points in the AEE are called *attacks* and represent a multitude of malicious actions, including port scanning and brute forcing. The AEE recorded more than 420 million attacks between January and December 2023. **That is more than 13 attacks per second, a 30% increase over what we observed in 2022.**

Our data is different from what is seen in many threat reports because it comes from specialized IT/OT/IoT honeypots that either *mimic* realistic device profiles – including exposed protocols, banners and parts of the filesystem – or *are* real specialized devices, instead of generic honeypots capturing every kind of attack.

A subset of these attacks contains *exploits* – attempts to exploit known vulnerabilities with a specific CVE identifier. **The intrusion detection systems connected to the AEE raised close to 320 million alerts related to vulnerability exploitation in the period of study.**

Our Malware Analysis Lab (MAL) collects and analyzes the malware samples dropped by attackers on the AEE or shared on public repositories. Our goal is not to analyze as many samples as possible but to focus on those that are unique. **The MAL has analyzed more than 50,000 unique malware samples between January and December 2023.**

Finally, we constantly hunt for new command and control (C2) infrastructure and maintain a threat actor knowledgebase with **data about more than 600 threat actors**.

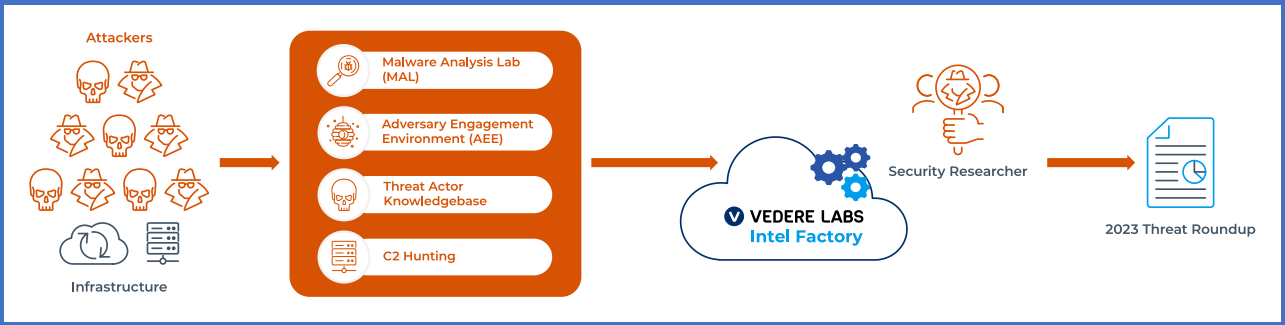

2. Main findings

2.1. The rise of China-based attacks

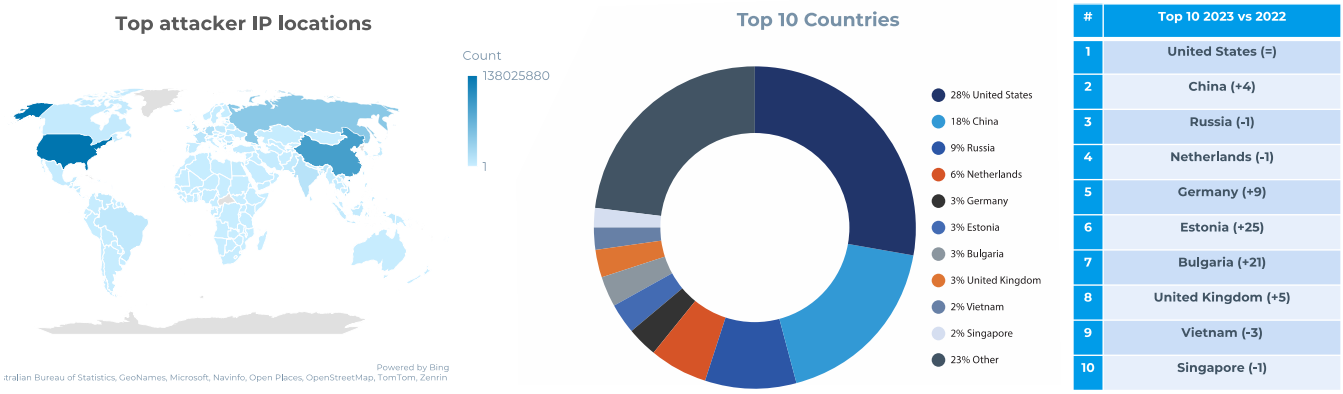

#	Top 10 2023 vs 2022
1	United States (=)
2	China (+4)
3	Russia (-1)
4	Netherlands (-1)
5	Germany (+9)
6	Estonia (+25)
7	Bulgaria (+21)
8	United Kingdom (+5)
9	Vietnam (-3)
10	Singapore (-1)

Figure 1 – Distribution of attacks by country

We detected attacks originating from 212 countries and territories (21 more than in 2022). Figure 1 shows the distribution of attacks detected by country of origin. Countries appear in this list due to the presence of legitimate hosting providers being abused by attackers; the presence of bulletproof hosting providers that cater specifically to cybercriminal activities; or the use of compromised hosts to launch attacks.

This year, the top 10 countries accounted for 77% of the malicious traffic, 4% more than in 2022. The top 10 list of countries originating attacks looks somewhat different from last year, with more European countries and a noticeable spike in attacks originating from China.

Insight for Defenders: Country of origin alone continues to be ineffective for judging the risk of a particular IP address because there are even more countries and territories originating attacks and because of the abuse of hosting providers and compromised devices in North American and European countries. The rise in attacks originating from China reinforces our 2022 message that if your organization does not do business with, or in, that country, blocking those IP ranges may help to reduce noise in the SOC.

2.2. Autonomous systems – more compromised devices

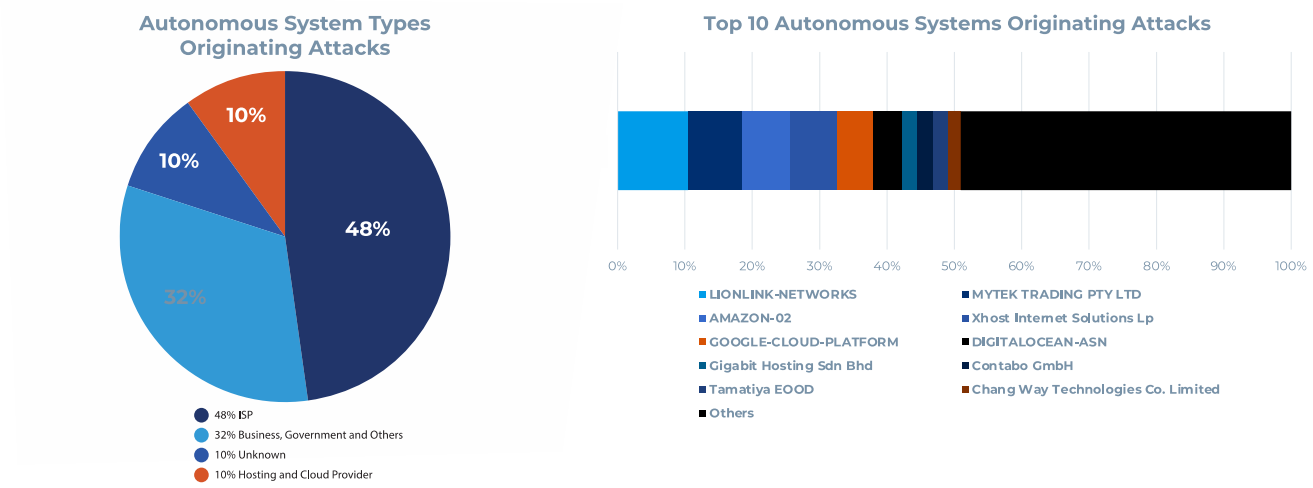

Figure 2 – Distribution of attacks by autonomous system

Attacks originated from more than 500 autonomous systems (ASes), which are blocks of IP addresses under the control of an organization. Figure 2 shows the percentage of attacks coming from the three types of ASes we observe:

- Internet Service Providers (ISPs) increased from 18% in 2022 to 48% in 2023.
- Business, Government, and others increased from 1% to 32%.
- Hosting or cloud providers decreased from 81% to 10%.

We changed our classification system for ASes in 2023, which explains some of the variation shown above. Nevertheless, the large chunk of attacks coming from ISPs as well as business, government and other organizations signifies an increase in the use of compromised devices to launch attacks instead of leasing infrastructure from dedicated providers. The variation could also be due to the increased popularity of "residential proxy" services, where threat actors proxy their traffic via applications running on residential devices, with IP addresses typically managed by ISPs. We discussed residential proxies in our 2023H1 Threat Review. On the cloud side, the use of Amazon and Google infrastructure for attacks has risen significantly, with those two alone responsible for more than 12% of attacks we observe.

Overall, the top 10 ASes are responsible for 52% of attacks, but only three ASes from the top 10 in 2022 remain in this year's list: LIONLINK-NETWORKS, Xhost Internet Solutions and – not surprisingly – DigitalOcean, which we reported last year as being one of the favorites for attackers and malware infrastructure.

Insight for Defenders: ASes are still a better sign of risk than country of origin. IPs belonging to known risky autonomous systems should always be treated with care. The increased use of residential proxies, ISPs and compromised devices on legitimate organizations means it is more important than ever to keep up-to-date with threat feeds that can monitor these compromised IP addresses and help to detect compromises in your own network.

2.3. Attacked services – focusing on the web

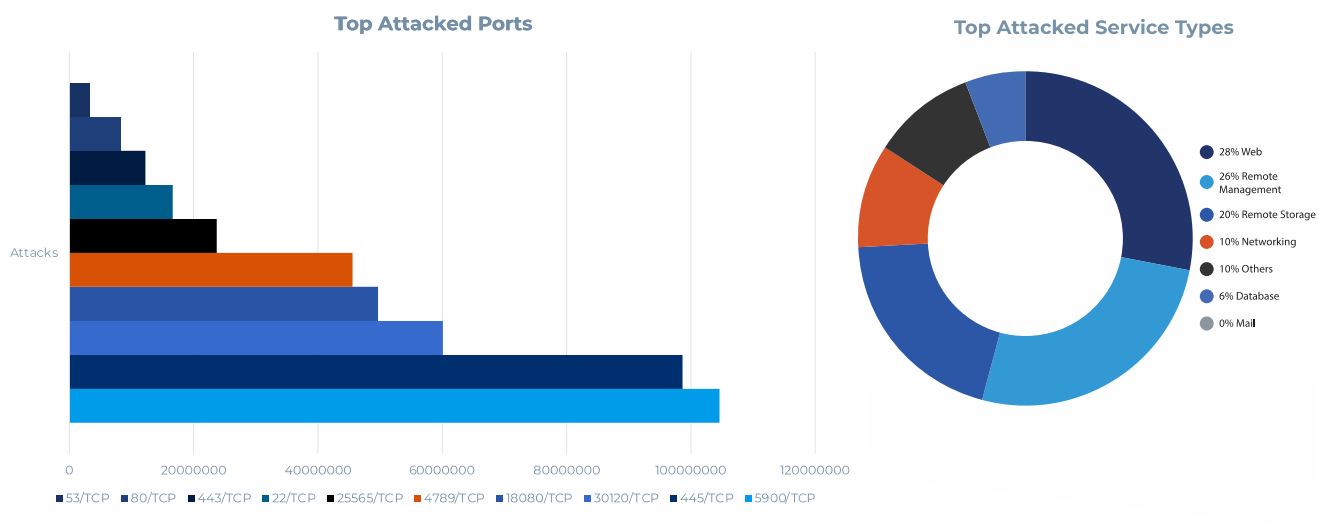

Figure 3 – Distribution of attacked ports and services

Figure 3 shows the share of traffic targeting each type of network service, classified according to assigned or well-known IPv4 TCP/UDP destination ports:

- Web applications increased from 26% in 2022 to 28% in 2023, becoming the most attacked service type. Most attacks against these services are either scanning or vulnerability exploitation attempts (see section 2.5).
- Remote management protocols, such as RDP and VNC for remote desktop and SSH and Telnet for remote terminal decreased from 43% last year to 26% this year. Attacks on these protocols are mainly brute forcing or password spraying (see section 2.4).
- Remote storage protocols, such as SMB and FTP, decreased from 23% to 20%.
- Networking protocols, such as DNS, DHCP and CWMP/TR-069, increased from 1% to 10%.
- Database services, such as Microsoft SQL Server, Redis, mongoDB, MySQL and PostgreSQL, increased from 1% to 6%.
- E-mail services, such as IMAP, POP3 and SMTP, remained unchanged at less than 1% of attacks.

Insight for Defenders: The increase in attacks targeting web and networking protocols is representative of threat actors shifting from mostly credential-based attacks to exploits on perimeter devices and applications. This was observed often in 2023 and means it is more important than ever to adopt technologies for risk management and threat detection that cover the entire attack surface, whether that is applications on a server, a perimeter device or an IT workstation.

2.4. Weak credentials – targeting IoT devices

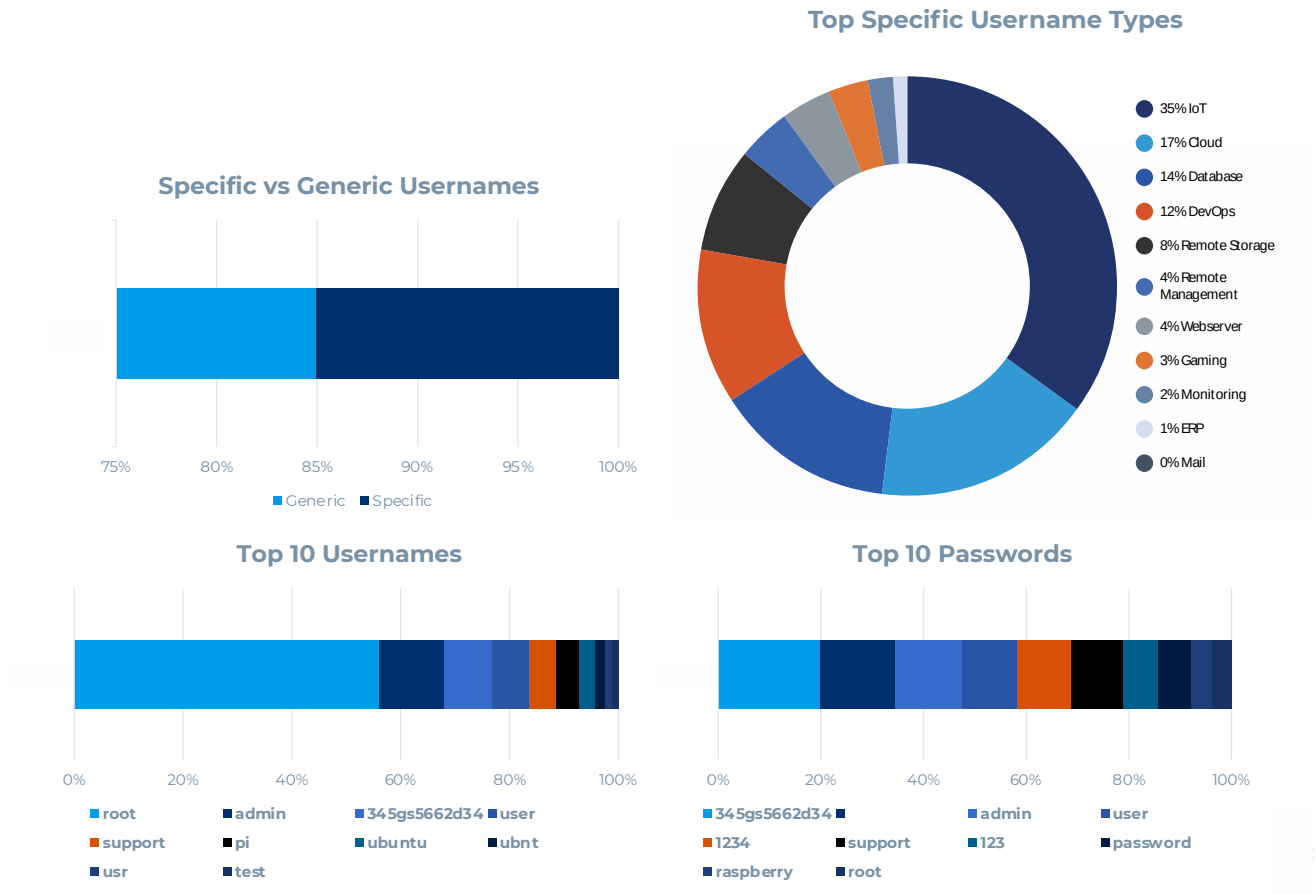

Figure 4 - Top abused credentials

Figure 4 shows the most abused credentials we observed, divided in two categories:

- Generic usernames (decreased from 87% to 85%) include “root,” “admin,” “user,” “guest” and several others.
- Specific usernames (increased from 13% to 15%) can be associated with specific roles, such as “www,” “backup,” “deployer” or even specific applications and devices, such as “odoo,” “rpi,” “kafka,” “zabbix” or “ec2-user”

The most interesting change we observed was the rise of specific usernames for IoT devices, which now account for 35% of this type of attack. These IoT-specific usernames include values such as “ubnt” (for Ubiquiti routers), “hikvision” (for IP cameras and DVRs), “moxa” (for industrial networking), “zyfwp” (for Zyxel firewalls) and “bigipuser3” (for F5 firewalls). Some of these are known backdoor accounts.

Insight for Defenders: Accounts for specific services are being scanned all the time, so make sure to change default usernames and passwords whenever possible. Try to use complex, unique passwords for every service on every device. Rotate credentials at a regular interval to avoid leaked credentials remaining valid. Finally, enable two-factor authentication whenever possible.

2.5. Exploits – there's much beyond KEV

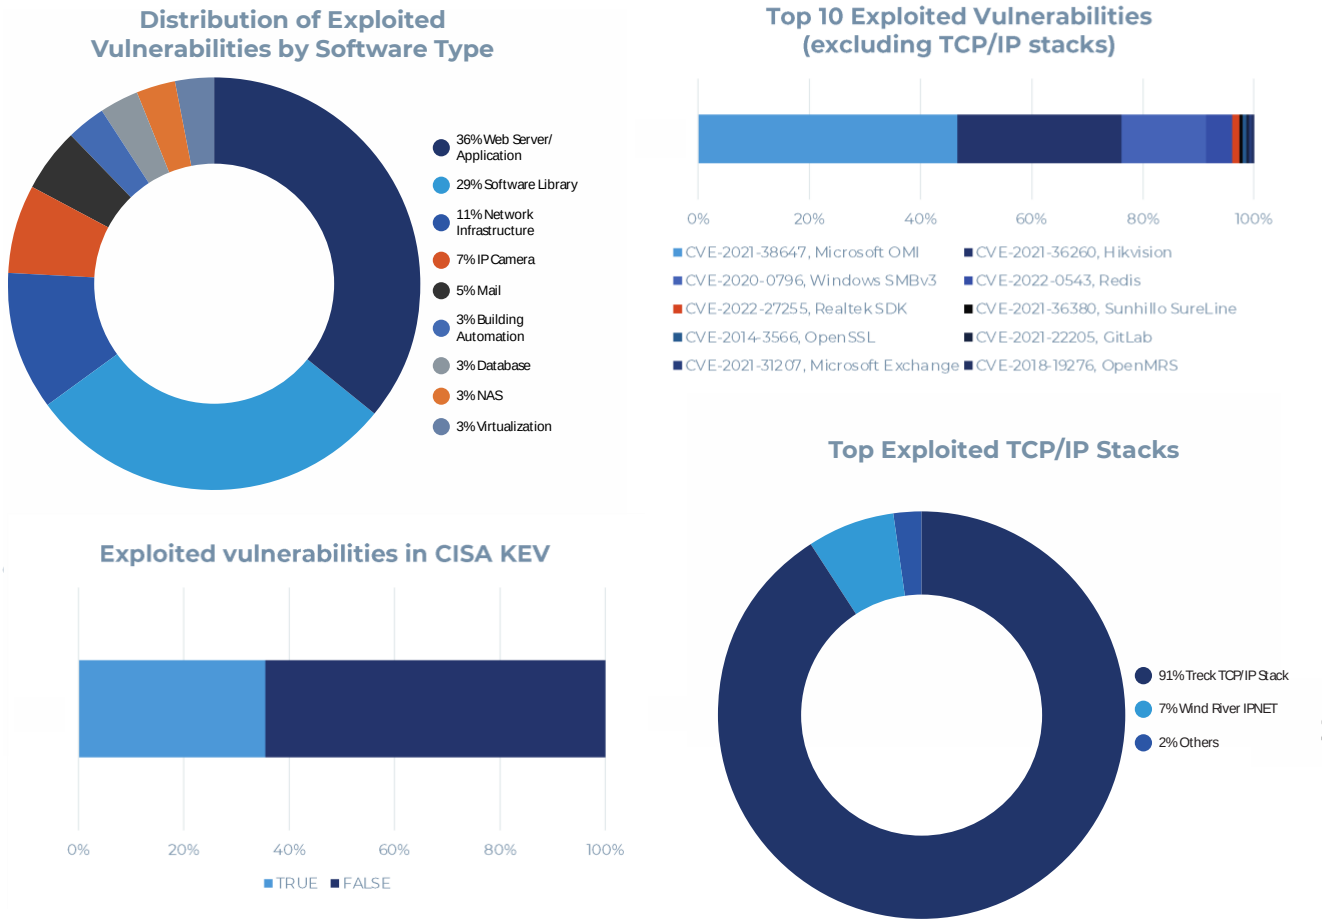

Figure 5 - CVEs exploited during the study period

Figure 5 shows the distribution of vulnerabilities we observed being exploited:

- Exploit attempts against web servers and applications rose from 14% in 2022 to 36% in 2023, becoming the largest category we see. This is in line with what we observed for targeted services in section 2.3.
- Software libraries (which include attacks against TCP/IP stacks) decreased from 76% to 29% of exploits. Part of this is explained by the significant decrease in massive exploitations of Log4j in 2023 compared to 2022.
- Exploits against network infrastructure devices, such as firewalls and routers, increased from 3% to 11%.
- Several categories of IoT devices known to be often exposed and vulnerable are massively targeted, such as IP cameras (7%), building automation (3%) and network attached storage (NAS, 3%).

We excluded exploits for TCP/IP stacks from the top 10 chart because detections for those can sometimes be noisy, leading to many similar vulnerabilities appearing in the top 10. If we look specifically at TCP/IP stack exploits, we observe that:

- Around 90% of those target the Treck TCP/IP stack, which was found vulnerable to Ripple20
- 7% target Wind River's IPNET, which is used in VxWorks and was found vulnerable to URGENT/11
- 2% of exploits target other TCP/IP stacks found vulnerable during Project Memoria, such as uIP and NicheStack

When looking at specific vulnerabilities being exploited, we notice that only 35% of those appear in CISA's Known Exploited Vulnerabilities (KEV) catalog.

Insight for Defenders: When deciding which vulnerabilities to patch and when, focus not only on CVSS and other severity metrics, but also consider the vulnerabilities that are currently being exploited. Although CISA keeps an up-to-date catalog of known exploited vulnerabilities, that does not cover the entirety of the exploited vulnerability landscape. Additional threat intelligence sources, **such as Vedere Labs' threat feeds,** are required to prioritize vulnerability risk mitigation.

2.6. OT attacks – a preference for popular protocols

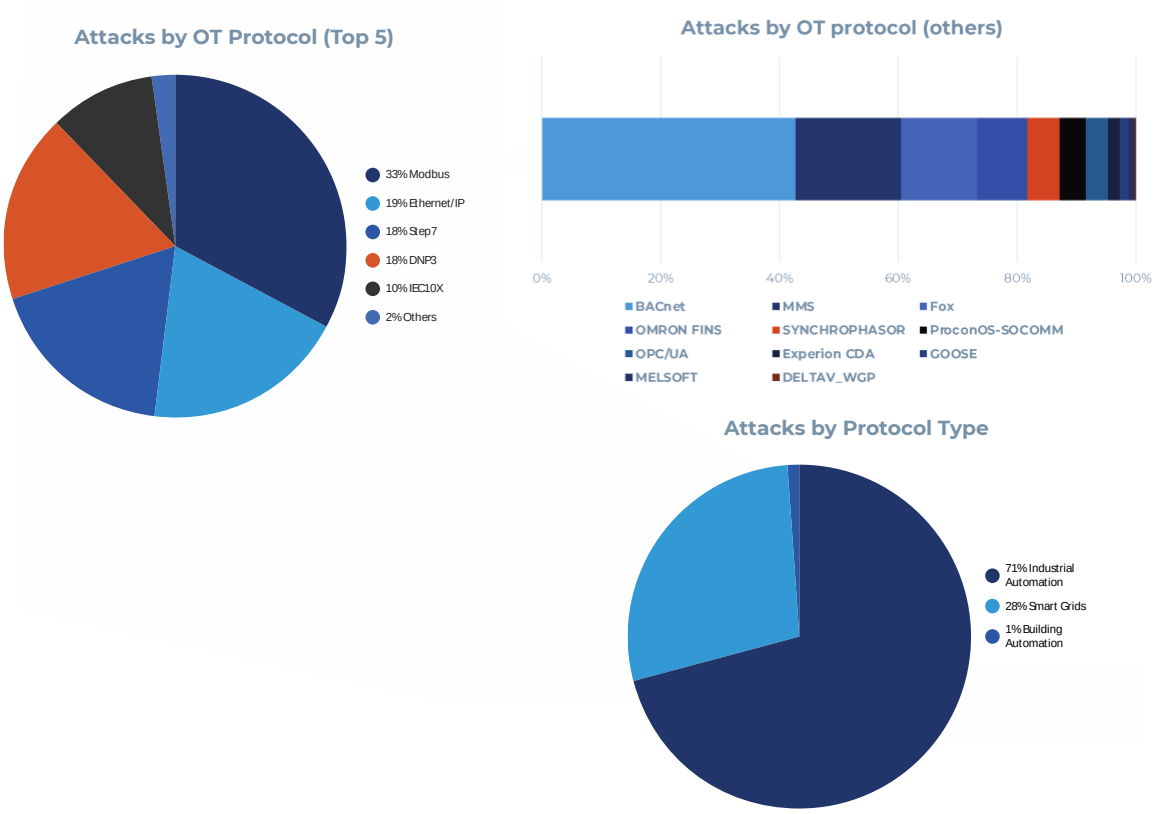

Figure 6 - Attacks against OT protocols

Figure 6 shows the distribution of attacks targeting OT protocols. There are five main protocols being constantly targeted: Modbus (a third of attacks), Ethernet/IP, Step7 and DNP3 (with around 18% each) and IEC10X with 10% of attacks. The remaining 2% are divided amongst many other protocols, of which the majority is BACnet.

As described last year, the majority of the activity in these protocols has to do with scanning and enumeration attempts. We also see messages with invalid, missing or truncated fields, which could cause devices to crash upon parsing them.

If we group these protocols into categories, we see that the majority of attacks target industrial automation protocols, such as Modbus, Ethernet/IP and Step7, which can be used in several critical infrastructure sectors. 28% of attacks focus on protocols used exclusively in the power sector, such as DNP3, IEC10X, MMS and GOOSE. Finally, around 1% of attacks focus on protocols used in building automation, such as BACnet and Fox. However, notice that *vulnerability exploits* against building automation devices are much more common than attacks leveraging their specific protocols, which is a different scenario than seen in the other automation categories (see section 2.5 for exploits against building automation).

Insight for Defenders: Monitoring the traffic to and from OT devices is nowadays as critical as monitoring IT traffic. Attackers are constantly probing these assets for weaknesses and many organizations will be blind to that because they lack visibility into their OT infrastructure. Building automation and even protocols such as Modbus are now found in almost every organization and are a target for attackers.

2.7. Attacker actions/TTPs – persistent threats

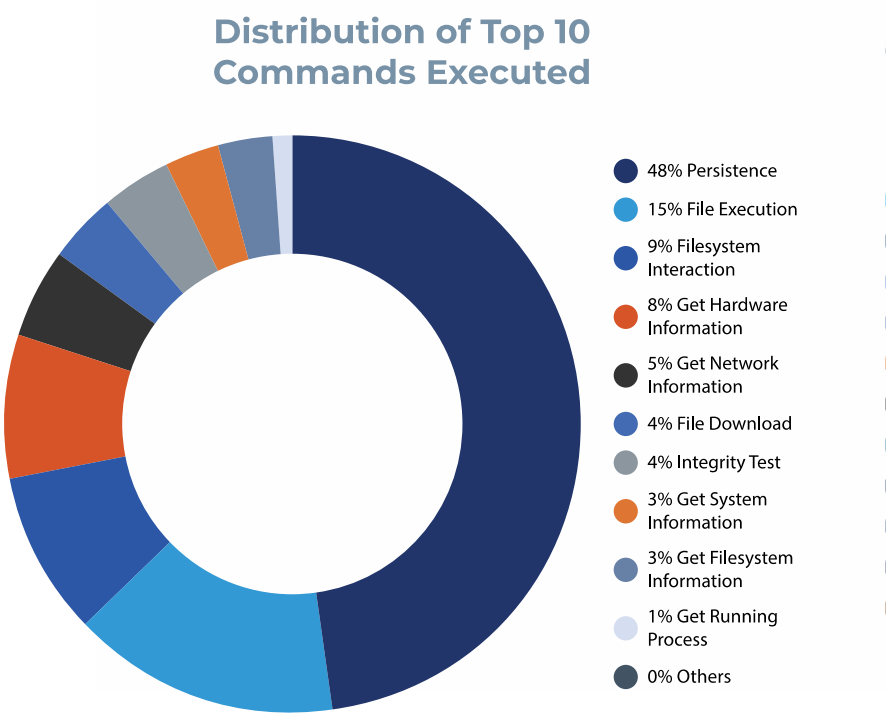

Figure 7 – Top executed commands

Figure 7 shows the distribution of top 10 commands executed, mainly over SSH or Telnet, after attackers managed to gain initial access. Most of the attacks we observed were automated and used the following ATT&CK tactics:

- TA0003 – Persistence represents around 50% of observed commands, up from 3% observed in 2022. Persistence comprises four main procedures: persisting SSH keys, downloading backdoored shells, creating or manipulating user accounts and executing background processes.
- TA0007 – Discovery represents around 25% of post-exploitation activities, down from 95% in 2022. These activities include obtaining information such as CPU, RAM, filesystem, operating system and architecture, as well as listing logged-in users, and running processes and scheduled jobs.
- TA0002 – Execution represents the other roughly 25% of observed commands, up from 1% in 2022. These commands are related to interacting with the filesystem, and downloading and executing additional malware.

Most of the commands observed target Linux shells, but we also saw actions specific to network operating systems, such as Cisco IOS (e.g., "show running-config"), Mikrotik RouterOS (e.g., "/ip cloud print" and others (e.g., "show mdm config," "show sip," "sip -print").

Insight for Defenders: The increase in persistence actions means that incidents are becoming harder to contain and eradicate after an initial breach. But it also means that threat actors intend to remain longer in a system, which reinforces our message last year that even after an initial breach, threat actors need to spend time getting situated in the target, downloading further tools, executing them and persisting. Many of these actions introduce more chances for detection and response, provided that proper endpoint inspection capabilities are available. Unfortunately, these capabilities are notoriously missing on non-IT endpoints.

2.8. Malware – RATs and infostealers

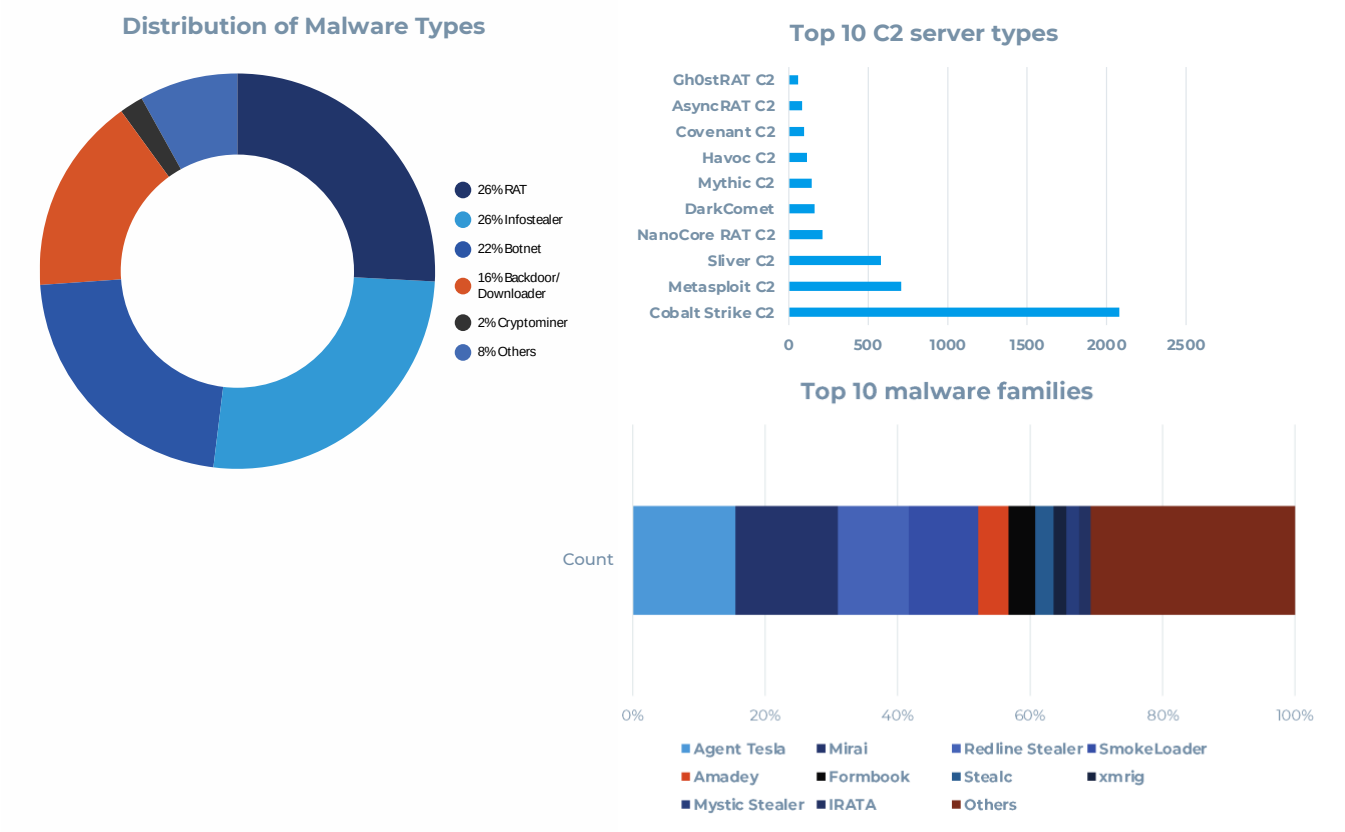

Figure 8 – Distribution of observed malware samples and C2 servers

Figure 8 shows the distribution of malware and observed command and control (C2) servers in our dataset. Last year we focused on botnets, but in 2023 we expanded our malware capture and observed an equal amount of remote access trojans (RATs) and information stealers (infostealers) as the most popular type of malware. Botnets and other downloaders come in third and fourth, respectively, followed by crypto miners and then a variety of other malware, such as keyloggers and adware.

The most popular individual malware families observed were Agent Tesla, a RAT that alone accounts for almost 16% of observations, then variants of the Mirai botnet (15%) and the Redline infostealer (10%).

Cobalt Strike remains by far the most popular C2 we see, with 46% of observations, followed by Metasploit (16%). The most interesting observation is the rising popularity of Sliver (13%), an open-source “adversary emulation framework” with a C2 and implants that support Windows, Linux and macOS. Most C2s are located in the United States (40%), followed by China (10%) and Russia (8%).

Insight for Defenders: Although individual malware samples and even families keep evolving every day, the basic nature of malware remains unchanged. The combination of RATs, botnets, infostealers and C2 servers is by now well-known to both attackers and defenders. As always, this means it is much more productive for defenders to detect and hunt for TTPs and anomalous behavior than to rely solely on file hashes and C2 IPs, which change constantly.

2.9. Threat actors – the reflex of geopolitics

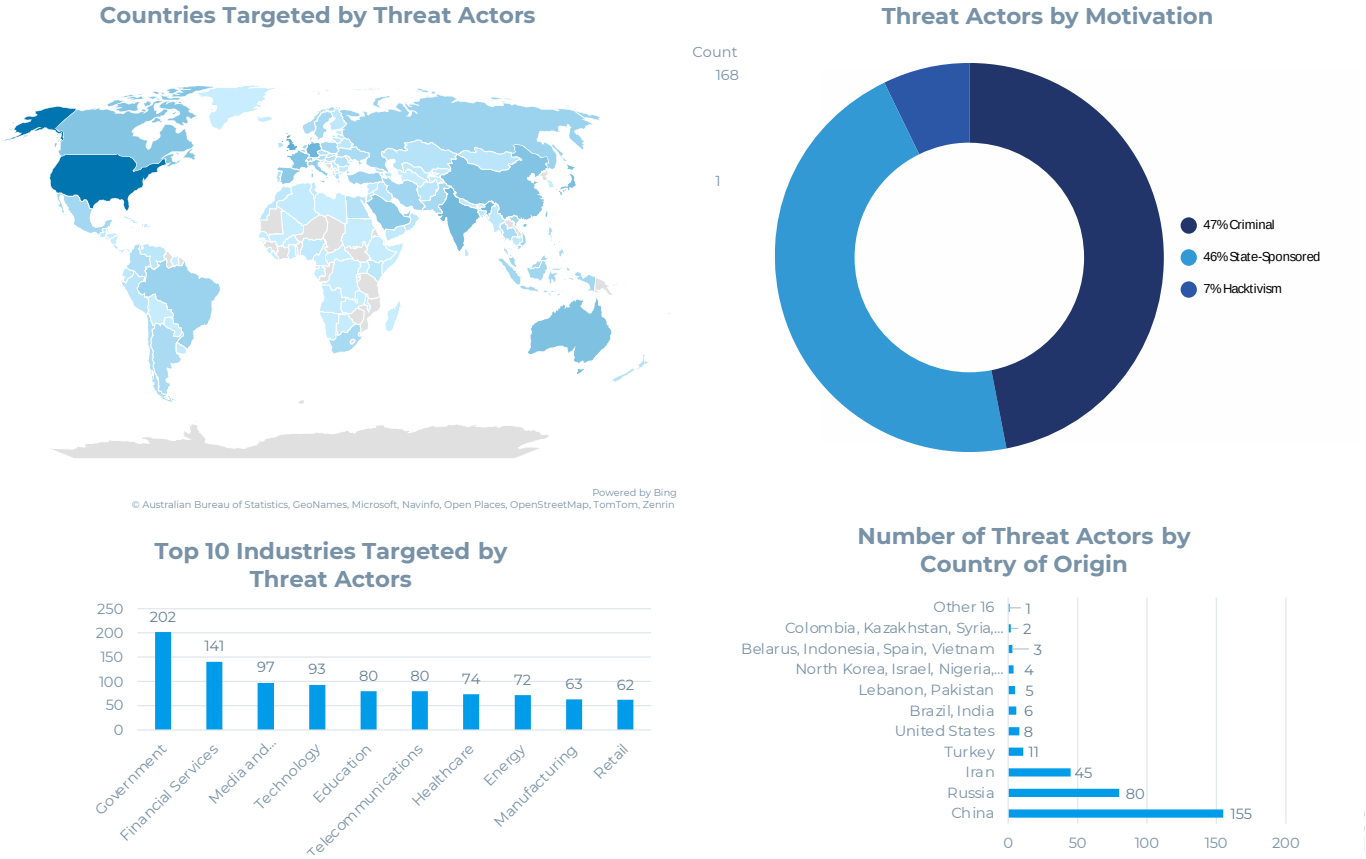

Figure 9 - Distribution of threat actors

We maintain a database of more than 600 threat actors, including cybercriminals, state-sponsored actors and hacktivists. Most threat actors we track are cybercriminals or state-sponsored, with a smaller number of hacktivists. Figure 9 shows how these actors were distributed in 2023:

- Threat actors have targeted 163 countries. The United States is the most targeted by far, with 168 actors aiming at the country. In second place comes the Unted Kingdom with 88, then Germany with 77, India with 72 and Japan with 66.
- Most actors were located in China (155), Russia (88) and Iran (45). Together, these three countries account for almost half of threat actor groups in our database.
- Government, Financial Services and Media and Entertainment were the industries most targeted by these actors.

Insight for Defenders: As we discussed in section 2.1, blocking communications simply by country of origin is not effective, but knowing where threat actors come from and their goals can help prioritize strategic security investments. Organizations in the most affected industries, especially, should pay attention to the latest threat intelligence to monitor campaigns that target specific sectors.

2.10 Targeted versus opportunistic attacks

Targeted attacks on Sierra Wireless routers

We observed close to 6,000 unique IP addresses attacking specialized Sierra Wireless AirLink OT/IoT routers deployed during our Sierra:21 research, both as software simulations and as real devices that attackers could fully interact with.

Several of those attacks targeted the devices specifically by exploiting CVE-2018-4068, CVE-2018-4070, CVE-2018-4071 and CVE-2018-4063. All four vulnerabilities were originally disclosed by Cisco Talos in 2019 and have public proof-of-concepts, which were only slightly modified by the attackers.

The first three vulnerabilities allow attackers to collect information about the targeted device, while the latter allows them to execute OS commands. All exploit attempts were followed by a successful login (we deliberately used weak credentials), since CVE-2018-4063 requires valid credentials. The end goal of these attacks was always to download malware onto the devices and enlist them into botnets.

Opportunistic attacks on every device

While studying attacks on those networking devices, we also noticed many reconnaissance attempts against *other* specific IoT and OT devices that happened to hit the Sierra Wireless routers. These attempts included fingerprinting for devices such as SonicWall firewalls, Siemens S7 PLCs, Tridium Niagara building automation controllers, Red Lion Controls HMI panels and VoIP phones from vendors such as Linksys, Mitel, ATCOM and Yeastar. They mostly used public NMAP scripts and came from commercial mass IPv4 scanners. However, even on the real devices, we also saw many other indiscriminate attacks, such as:

- Web application/server exploits for vulnerabilities such as CVE-2013-6397, CTX129430, CVE-2019-18935, CVE-2019-9670, CVE-2017-9841, CVE-2021-26086, CVE-2021-44228 (log4j) and others.
- Attempts to obtain information from misconfigured web applications by grabbing sensitive files, private keys or database endpoints.
- Malware deployment via OS command injection vulnerabilities in several networking devices. These exploits were mostly used to drop variants of Mirai, Gh0sT RAT and SystemBC C2 communications.

Insight for Defenders: Specialized devices are being exploited both in targeted and opportunistic attacks simultaneously. Obviously, targeted attacks are more worrying because the adversaries know what they are looking for, but some opportunistic attacks can also reveal more information that attackers will use to expand their attack campaigns. Continuously identify and patch vulnerable devices and segment networks to ensure that “low hanging fruit” such as known vulnerable edge devices cannot lead to further compromises on your network.

Conclusion

In this report, we analyzed data about attacks, exploits and malware we observed in 2023. Throughout, we included insights for defenders alongside each of the main findings. At a more strategic level, we recommend organizations focus on three key pillars of cybersecurity**:**

- **Risk & Exposure Management.** Start by identifying every asset connected to the network and its security posture, including known vulnerabilities, credentials and open ports. Then, change the default "easily guessable" credentials and use strong, unique passwords for each device. Next, unused services should be disabled and vulnerabilities patched to prevent exploitation. With your attack surface understood, you can now fully assess risk in your environment. Finally, focus on mitigating using a risk-based approach. Use automated controls that do not rely only on security agents and apply to the whole enterprise instead of silos like the IT network, the OT network or specific types of IoT devices.

- **Network Security.** Do not expose unmanaged devices directly on the internet. Segment the network to isolate IT, IoT and OT devices, limiting network connections to only specifically allowed management and engineering workstations or among unmanaged devices that need to communicate. Segmentation should not happen only between IT and OT, but even *within* IT and OT networks to prevent lateral movement and data exfiltration. Restrict external communication paths and isolate or contain vulnerable devices in zones as a mitigating control if they cannot be patched or until they can be patched.

- **Threat Detection & Response.** Use an IoT/OT-aware, DPI-capable monitoring solution to alert on malicious indicators and behaviors, watching internal systems and communications for known hostile actions such as vulnerability exploitation, password guessing and unauthorized use of OT protocols. Anomalous and malformed traffic should be blocked, or at least alert network operators to its presence. Beyond network monitoring, extended detection and response (XDR) solutions are an important consideration. They collect telemetry and logs from a wide range of sources, including security tools, applications, infrastructure, cloud and other enrichment sources, correlate attack signals to generate high-fidelity threats for analyst investigation and provide the ability to automate response actions across the enterprise.

The most important takeaway is that the traditional cyber hygiene practices mentioned above must address *every* asset on the network, prioritizing the most critical attack surface based on up-to-date threat and business intelligence.