FORESCOUT RESEARCH | VEDERE LABS

Exposing the Exploited

A quantitative analysis of vulnerabilities under the radar

May 7, 2024

Contents

Our data source: The VL-KEV

Forescout Research – Vedere Labs' catalog of known exploited vulnerabilities (VL-KEV) is a list of vulnerabilities that we observe exploited in the wild from two sources:

1. The Adversary Engagement Environment (AEE) is a set of honeypots on the open internet luring attackers and recording their actions. These specialized IT/OT/IoT honeypots 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.

2. A threat actor knowledgebase with data about more than 600 threat actors coming from internal and third-party reports. The data includes vulnerabilities that have been observed exploited by the tracked actors.

1. Executive summary

There are too many gaps in the named security vulnerability process. And there are plenty of vulnerabilities that do not receive the attention they deserve. Some vendors silently patch issues while others leave vulnerabilities in a reserved state. There is not one source of information that contains every vulnerability being exploited.

The result? Major gaps in time and in your security team's responsiveness and effectiveness.

Our latest research examines this world of exploited vulnerabilities outside of standard catalog systems. This report provides a real picture of the vulnerabilities exploited in the wild to see beyond the hype of mass-exploited vulnerabilities – with a focus on vulnerabilities rarely discussed.

Unfortunately, organizations are too reliant on these vulnerability catalogs, such as CISA's Known Exploited Vulnerabilities (KEV), FIRST's Common Vulnerability Scoring System (CVSS) scores and others. Each official vulnerability data source is crucial, yet each has limitations.

Why? Because the definition of what is exploited varies among these official data sources. Some consider only what is currently exploited. Others consider a vulnerability to be exploited only once it is in the wild. Vulnerabilities without a CVE identifier and CVSS scores that do not reflect real risk are well-known in the security industry. Exploit predictions are useful to understand what *may* happen in the future.

But what about currently exploited vulnerabilities in the wild that are *not officially recognized yet*?

The bottom line: Defenders need other sources of threat intelligence to help manage the volume and frequency of exploits. Only paying attention to named catalogs is dangerous as vulnerabilities are being discovered, weaponized and exploited in the wild faster than ever before. With 31 zero days already in 2024, it is time to take stock of the challenge.

Key findings:

- Nearly 90,000 vulnerabilities are without a CVE ID
- Thousands of devices are affected by 28 vulnerabilities in our catalog (untracked by CISA)
- 83% of exploited vulnerabilities have either high or critical CVSS scores
- More than 21,200 issues were discovered in 2023 with an unassigned CVE ID
    - Up 4% from 2022
    - And up 45% from 2021
- 44% of the vulnerabilities without a CVE ID can be used to gain access to a system
    - 37% have high or critical severity
    - 45 exploited vulnerabilities did not have a CVE ID (2.15% of the total)
- A total of 2,087 distinct exploited vulnerabilities seen across four databases:
    - CISA KEV, AttackerKB, Shadowserver, VL-KEV
    - No database alone contained all the information
    - CISA had 50% of the total exploited vulnerabilities (1055)
    - 47% are seen in only one database (968)
    - Only 4% are seen in all four (90)
- The most exploited OT and IoT devices are:
    - Network Attached Storage (NAS)
    - IP cameras
    - Building automation devices
    - VoIP equipment

2. Research findings

We use CISA KEV criteria to guide our research into the current vulnerability landscape and help define what is exploited. Our focus is on **what** is exploited and **how**. And we also explore vulnerabilities without CVE IDs and offer remediation guidance.

1. Vulnerabilities without a CVE ID

KEY FINDING	WHY IT MATTERS	REMEDIATION
There are 62% more vulnerabilities *without* a CVE ID in 2023 than in 2021.	Orgs that rely exclusively on vulnerabilities with CVE IDs are *blind* to thousands of issues – with more on the way.	Include multiple sources of vulnerability information and threat intelligence *outside* of CVEs in the NVD.

The CVE system has more than 26,000 issues identified in 2023. But it is missing a lot more vulnerabilities than you might expect on devices used globally. The prevalence of vulnerabilities outside the CVE ecosystem is driven by modern issues, such as a fragmented open-source landscape. This is evidenced by the fact that the number of vulnerabilities without a CVE ID is constantly growing.

Alternative catalogs, such as the Chinese vulnerability databases CNVD and CNNVD, include anywhere between 1,600 and 12,000 *more* vulnerabilities than the US NVD. This is due to a very different ecosystem where researchers and vendors must notify the Chinese government of vulnerabilities found even before patches are made available.

It is easy to find examples of vulnerabilities affecting Chinese devices commonly used in the West that do not have a CVE ID. For example, there were 64,125 Chinese-made Ruijie routers exposed on the Internet at the time of writing this report. **Some of these routers are vulnerable to at least two issues that we see exploited but have no CVE IDs: CNVD-2021-09650 and SSV-89107.** They are most popular in China (73%), but there are thousands in the US (6,155) and UK (1,197), as well as hundreds in Japan (839) and several European countries, including Germany (790).

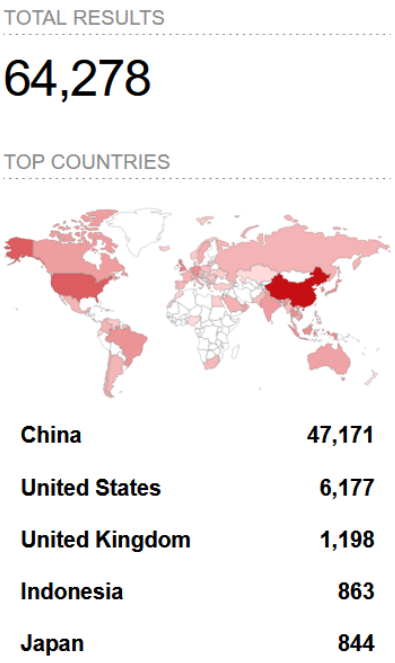

Figure 1 – Ruijie routers around the world

But this isn't just an issue of country of origin. There are vendors that silently patch vulnerabilities or prefer not to assign CVE IDs to issues discovered on their products. Plus, there are vendors that take a very long time to publish information about vulnerabilities — and leave the CVEs in a "RESERVED" state in the meantime [see our deep dive on the topic: OT:ICEFALL research]. Plus, there can be delays on NVD analysis of vulnerabilities.

We analyzed three current databases that aggregate information from several sources of vulnerabilities to provide a comprehensive picture. Table 1 lists these three databases.

Figure 2 shows that all three include between 7% and 29% of vulnerabilities without a CVE ID; This equates to between 21,000 and 40,000 vulnerabilities without a CVE ID per database; And up to 90,000 in total if they are all distinct issues.

Table 1 – Vulnerability databases

Dataset	Description	Total Vulnerabilities (Dec. 2023)	Vulnerabilities without a CVE ID (Dec. 2023)
Open-Source Vulnerabilities (OSV)	An open database launched by Google in 2021 with the goal of tracking vulnerabilities in open-source components using the OSV schema.	98,196	28,401
Global Security Database (GSD)	An open database launched by the Cloud Security Alliance (CSA) in 2022 with the goals of replacing CVEs and including new types of vulnerabilities, such as those in smart contracts.	317,555	21,521

IBM X-Force Exchange	A database maintained by IBM aggregating vulnerabilities in the NVD and several other sources.	248,826	40,304

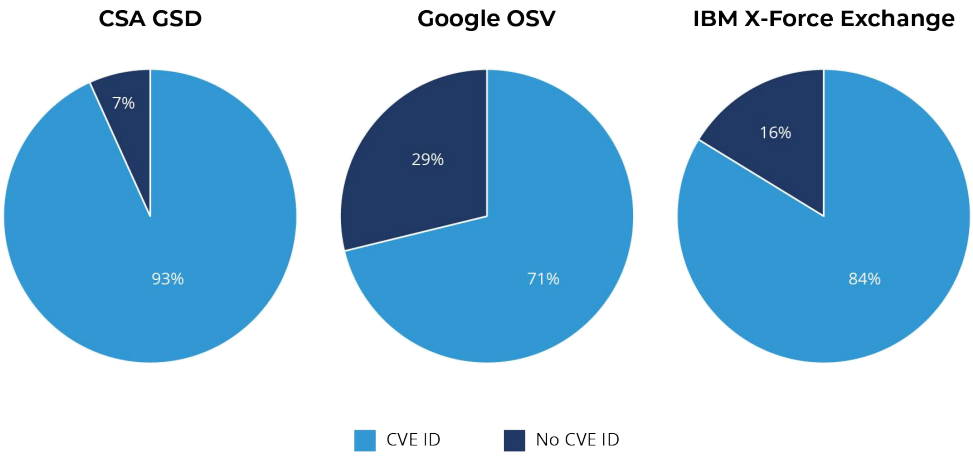

Figure 2 – Distribution of vulnerabilities without a CVE ID in each database

Figure 3 shows that, there were more than 21,200 issues discovered in 2023 and not assigned a CVE ID; which is 4% more than the 20,400 in 2022; which is 45% more than the 11,200 in 2021.

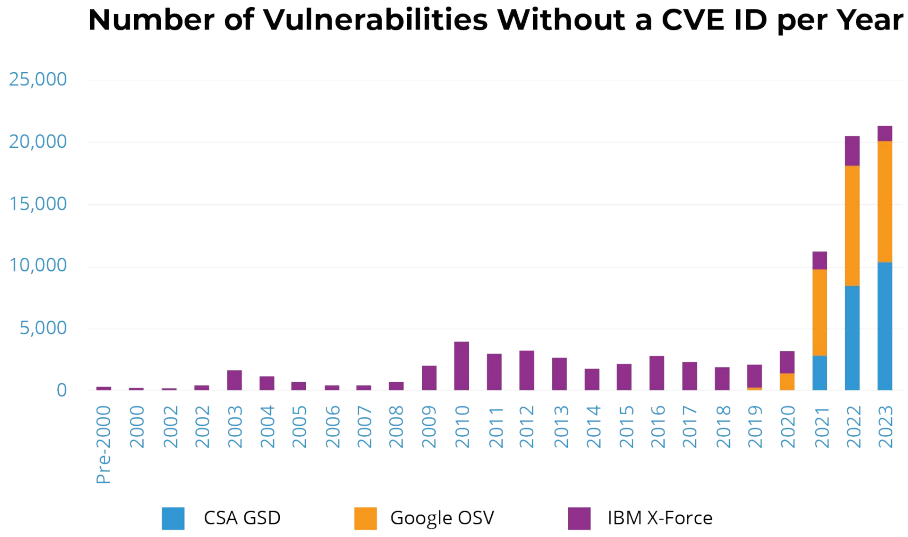

Figure 3 – Number of vulnerabilities without a CVE ID per year

There is no consistent way to score vulnerabilities across databases – which is done by CVSS scores in the CVE ecosystem. IBM X-force has details about the consequence and risk for the majority of vulnerabilities in their database, even those without a CVE ID. Figure 4 shows that 44% of the vulnerabilities without a CVE ID can be used to gain access to a system, which is the most common consequence for those issues. At the same time, 56% of these vulnerabilities are considered medium risk (with a score between 4.0 and 6.9) — while 37% have either high or critical severity.

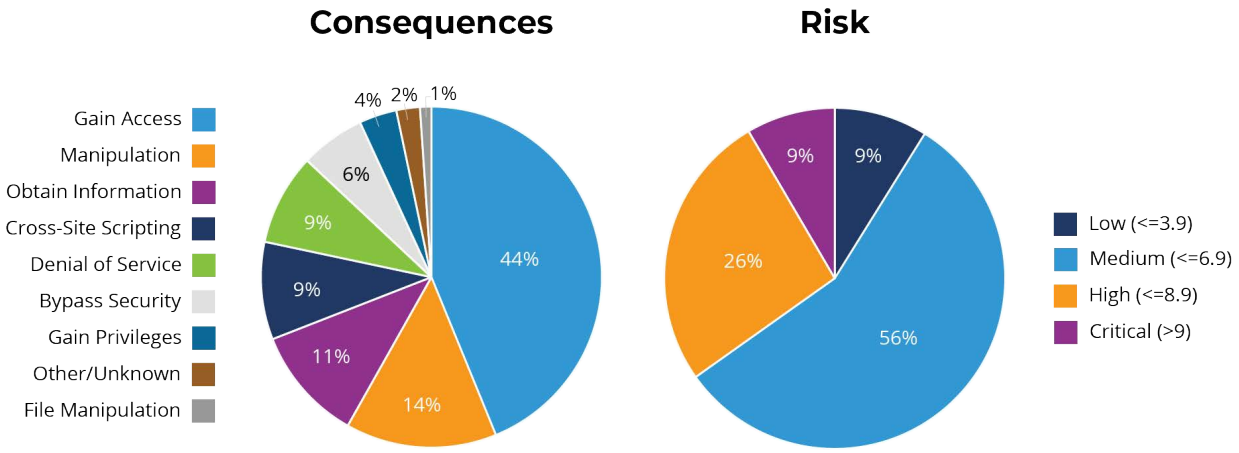

Figure 4 – Distribution of consequence and risk of vulnerabilities without a CVE ID on IBM X-force

2. Exploited vulnerabilities

The definition of what is exploited can change depending on the information source. For instance, once a vulnerability is added to the CISA KEV, it is not removed afterwards. Whereas sources based on honeypot data tend to be 'dynamic' in the sense that they reflect vulnerabilities that are *currently* being exploited. Some lists are manually updated based on human analysis. Others are automatically updated. For example, they could be based on detection rules from monitoring systems.

To understand what vulnerabilities should be considered 'exploited', we use four different databases [Table 2].

Table 2 – Databases of exploited vulnerabilities

Dataset	Description	Total Exploited Vulnerabilities (Dec. 2023)	Unique Exploited Vulnerabilities (Dec. 2023)	Data source type
CISA KEV	An open catalog maintained by CISA and currently the most popular source of information about exploited vulnerabilities. Criteria for inclusion are discussed at the beginning of Section 2.	1,055	31	Manually updated
AttackerKB	A forum discussing vulnerability exploitation maintained by Rapid7. The information is provided by the community (crowdsourced) and researchers can tag vulnerabilities as exploited in the wild. It includes most of CISA KEV as well as other vulnerabilities spotted by the community.	1,460	381	Manually updated

Shadowserver	A list of exploited vulnerabilities captured by Shadowserver's honeypots.	572	357	Dynamic, automatically updated
VL-KEV	A database maintained by Forescout Vedere Labs including vulnerabilities observed in our honeypots (AEE) or reported to be used by threat actors in our knowledgebase.	596	199	Partially dynamic, automatically updated

1. Database uniqueness

KEY FINDING	WHY IT MATTERS	REMEDIATION
47% of exploited vulnerabilities are tracked in only one database with the remaining 53% in multiple databases.	Relying on a single source of exploit information leads to ignorance hundreds of exploited vulnerabilities.	Use multiple sources on exploited vulnerabilities to prioritize patching.

There was a total of 2,087 distinct exploited vulnerabilities seen across the databases, but no database alone contains all the information. The database with **the most vulnerabilities was AttackerKB**, with 1,460 (70% of the total), but it relies on community information that potentially includes false positives. The database with **the least vulnerabilities was Shadowserver**, but that is because it relies solely on honeypot data which will miss the actions of targeted APTs, for instance and only includes timely information, not historically exploited issues. **CISA had 1,055 or 50% of the total exploited vulnerabilities**.

Each database contained several "unique" vulnerabilities which are those that are only reported by that one database and no other. The database with the most unique vulnerabilities was AttackerKB. The one with the least unique vulnerabilities was CISA KEV because most of their information is included in AttackerKB by the community.

Overall, 968 exploited vulnerabilities (47% of the total) are seen in only one database. Only 90 (4%) are seen in all four, as shown in Figure 5. **That means that relying on any one database alone can be dangerous**.

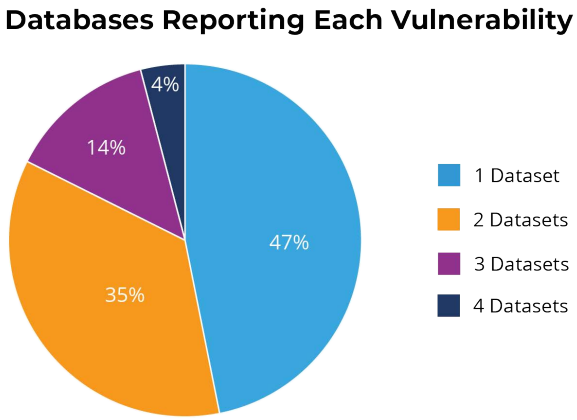

Figure 5 – Number of datasets where each vulnerability is reported as exploited

To understand the importance of having more coverage of exploited vulnerabilities, we analyzed Forescout Risk and Exposure Management customer data and saw 8,162 devices affected by 28 vulnerabilities in VL-KEV that are not on CISA KEV. The most impacted devices were distributed as shown in Figure 6, with UPSs, computers and printers accounting for around 25% each, infusion pumps and network equipment following with 9% and 8%, respectively, and then several other types of IoT, OT and IoMT devices with between 1% and 3%. 308 of those devices were exposed to the internet.

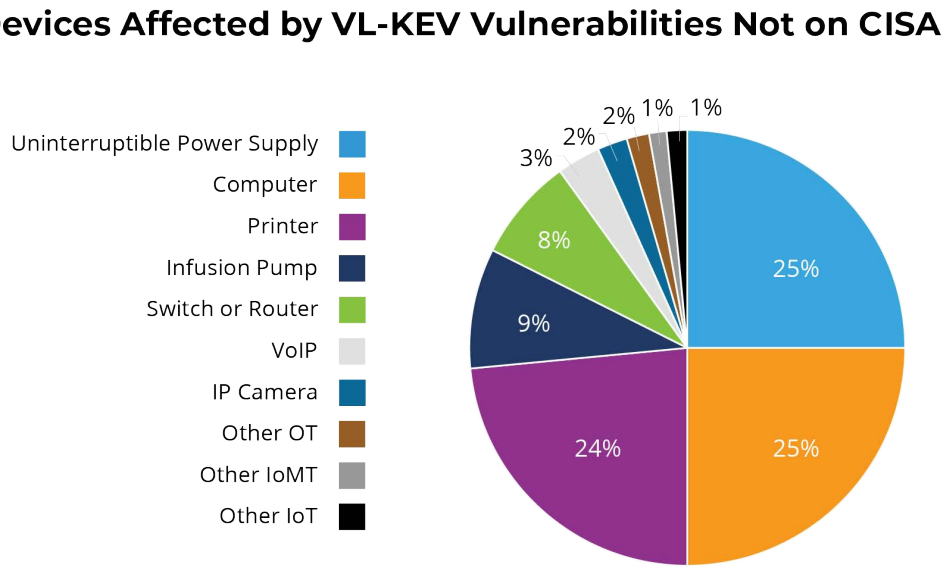

Figure 6 – Devices on customer networks affected by VL-KEV vulnerabilities not on CISA

2. Severity and root causes

KEY FINDING	WHY IT MATTERS	REMEDIATION
83% of exploited vulnerabilities have either high or critical CVSS scores	Web application security is still being ignored and easily exploited.	Add extra layers of protection for web applications including web app firewalls.

Forty-five (45) of the exploited vulnerabilities did not have a CVE ID. That includes 28 vulnerabilities with an ExploitDB (EDB) ID, 11 with a CNVD ID and 6 with other types of identifiers.

For the exploited vulnerabilities with CVE IDs, Figure 7 shows that most had either high (44%) or critical (39%) severity. The most common CVSS score was 9.8 for 571 vulnerabilities — while only 92 had a score of 10.

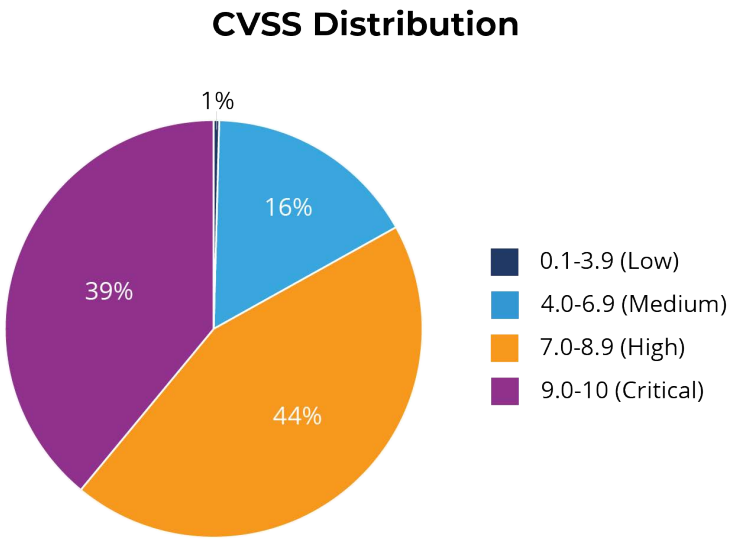

Figure 7 – Distribution of exploited vulnerabilities per CVSS score

The most common root causes for these exploited issues categorized by Common Weakness Enumerations (CWEs) are shown in Figure 8. The most common CWEs were divided into three 'types':

- Web applications – such as, command injections, path traversals and cross-site scripting
- Memory management issues – such as, out-of-bounds write and use after free
- Others

The most common CWE was "OS command injection" which was the root cause of 160 exploited vulnerabilities. Overall, CWEs related to web application issues were the most common. Beyond the ten root causes shown in Figure 8, there were 126 others that applied to 898 vulnerabilities.

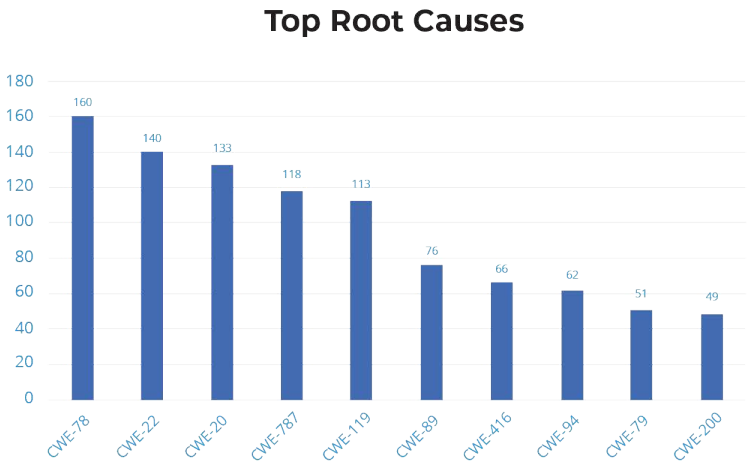

CWE	TITLE	TYPE
CWE-78	OS Command Injection	Web
CWE-22	Path Traversal	Web
CWE-20	Improper Input Validation	Web
CWE-787	Out-of-Bounds Write	Memory
CWE-119	Improper Restriction of Operations Within the Bounds of a Memory Buffer	Memory
CWE-89	SQL Injection	Web
CWE-416	Use After Free	Memory
CWE-94	Code Injection	Web
CWE-79	Cross-Site Scripting	Web
CWE-200	Exposure of Sensitive Information to an Unauthorized Actor	Other

Figure 8 – Top root causes of exploited vulnerabilities

3. Vulnerability age

KEY FINDING	WHY IT MATTERS	REMEDIATION
61% of exploited vulnerabilities across multiple sources were disclosed since 2020.	More recent vulnerabilities are more likely to be exploited, so there is often little time to patch before exploitation.	Focus on mitigating more recent vulnerabilities though understand there can be blind spots.

Overall, more recent vulnerabilities are more likely to be exploited. In 2020, there was a sharp increase. Based on all available data, 55% of exploited vulnerabilities have been disclosed since 2020. Honeypot data (Shadowserver and part of VL-KEV) shows 61% of vulnerabilities were disclosed since 2020. Figure 9 shows the distribution of exploited vulnerabilities by year of disclosure.

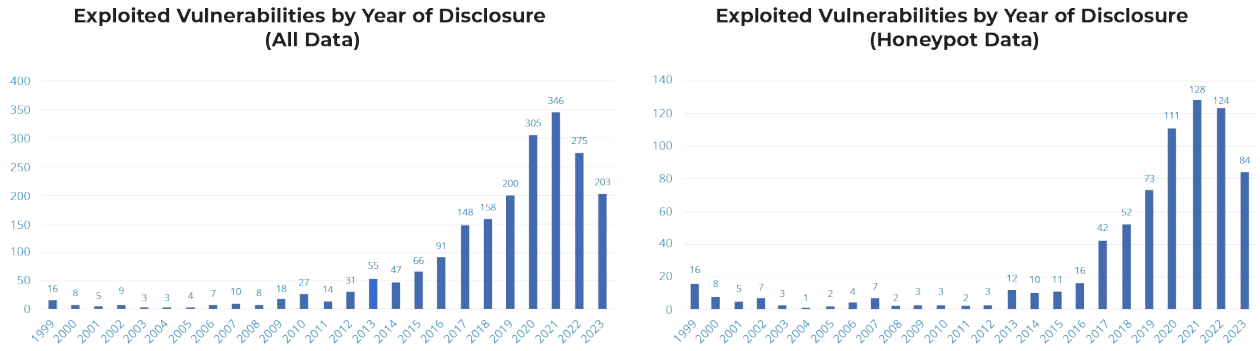

Figure 9 – Distribution of exploited vulnerabilities per year

4. Exploited targets

KEY FINDING	WHY IT MATTERS	REMEDIATION
Most exploited vulnerabilities affect web applications, operating systems, network devices, desktop software and OT/IoT devices.	OT/IoT devices are unmanaged and riddled with exploited vulnerabilities and unavailable patches which means they may stay unpatched for longer and become easier targets.	Prioritize those assets with internet exposure.

The most common targets were:

- Web applications: 22%
- Operating systems including TCP/IP stacks: 18%
- Routers: 12%

Most exploited OT and IoT devices:

- Network Attached Storage (NAS): 21%
- IP cameras: 17%
- Building automation devices: 12%
- VoIP equipment: 11%

Figure 10 shows the most exploited product types.

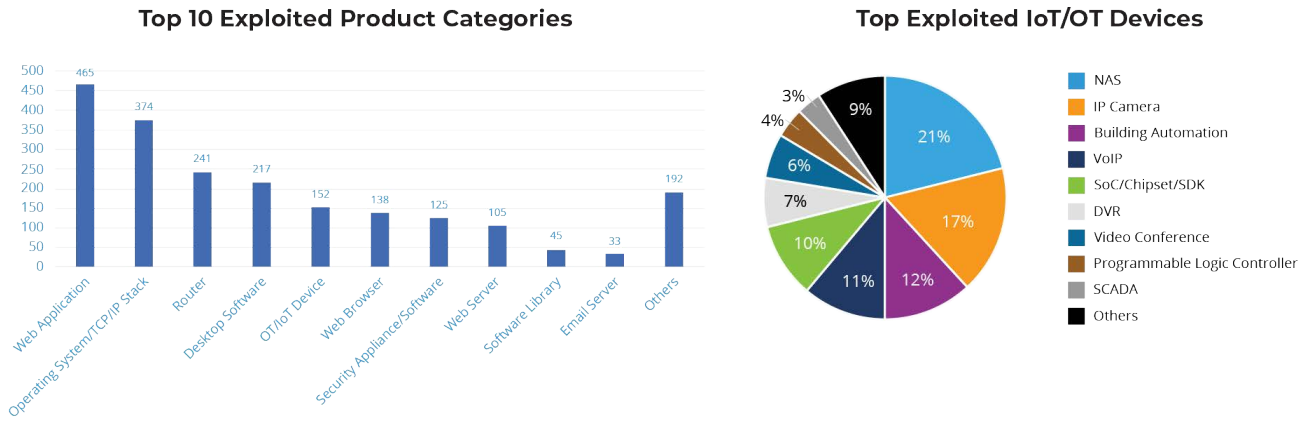

Figure 10 – Distribution of exploited vulnerabilities per product and device type

5. Exploit payloads

KEY FINDING	WHY IT MATTERS	REMEDIATION
50% of vulnerabilities were exploited with between two and 10 different payloads and for a short period of time.	Detection signatures may miss payload variants, especially when they change quickly.	Use network and host detection systems and ensure that their signatures are updated. Focus on network detection for unmanaged devices, such as OT/IoT.

After understanding *what* is exploited, we focused on *how* these vulnerabilities are exploited. To do so, we used only the data from VL-KEV vulnerabilities coming from AEE, since those are the only for which we have information about exploit payloads and who was exploiting them.

Figure 11 shows how four characteristics of the exploited vulnerabilities are distributed:

- **Occurrences:** How many times we observed the vulnerability being exploited. For 13% of vulnerabilities, we observed a single exploitation attempt with a payload. For 38%, we observed between two and ten attempts. For 24%, we observed between 11 and 100. That means that three quarters of the vulnerabilities were exploited up to 100 times. Only 1% of vulnerabilities were exploited more than 10,000 times.
- **IPs:** How many different IP addresses attempted to exploit the vulnerability. One third of vulnerabilities were always exploited by the same IP address, while almost half were exploited by between two and ten IPs.
- **Days:** The difference in days between the first observed exploitation of the vulnerability and the last observed exploitation. 16% of vulnerabilities were exploited on a single day, while more than half (53%) were exploited for between two and ten days. Only 13% of vulnerabilities were exploited for more than 100 days.
- **Payloads: The number of different ways we saw the vulnerability exploited.** Here, we use a strict definition of different: Even a single character difference in a payload that would achieve the same result counts as a different payload.
    - More than a quarter of vulnerabilities were always exploited with the exact same payload.
    - **Half of them were exploited with between two and ten different payloads**.
    - Only 2% of vulnerabilities were exploited with more than 1,000 different payloads.
    - In most cases, these payloads come directly from or are small changes to public proof-of-concept (PoC) exploits.

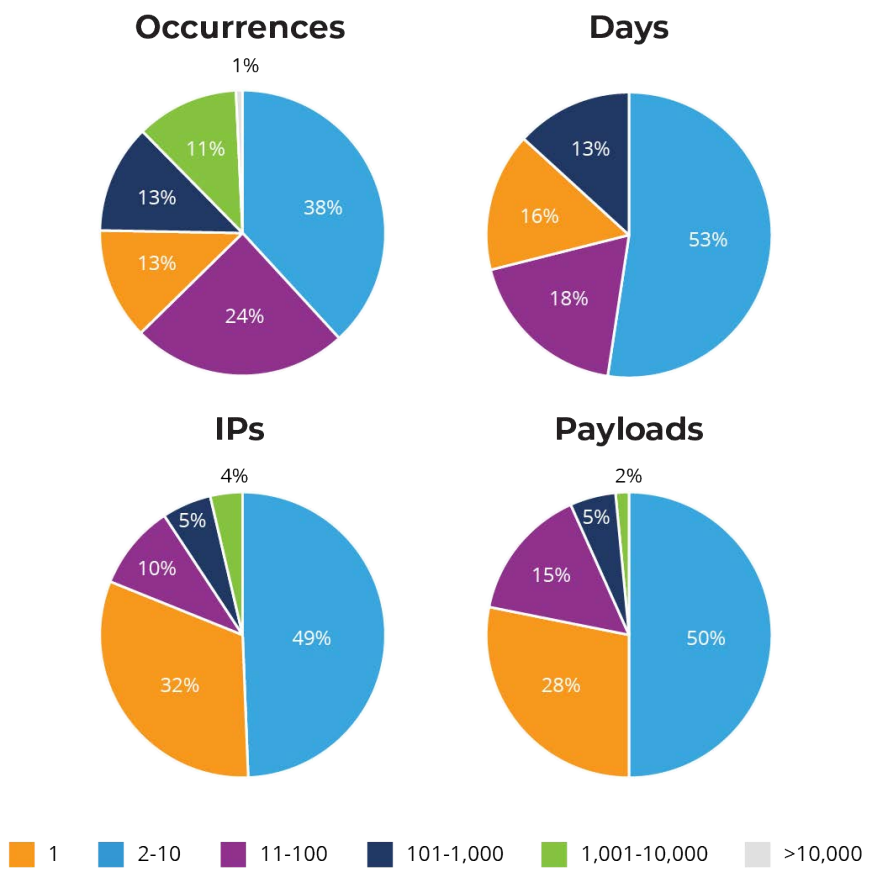

Figure 11 – Distribution of exploited vulnerabilities per number of occurrences, number of distinct payloads, number of attacker IPs and days when it was exploited

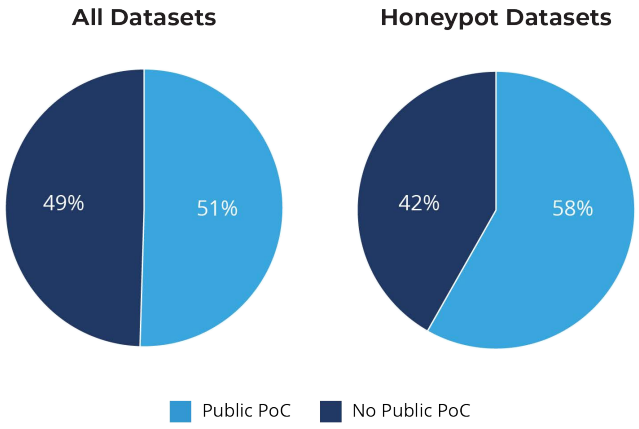

Figure 12 – Distribution of exploited vulnerabilities with and without a public PoC

Figure 12 shows that we could find – via a quick automated search on ExploitDB and GitHub – public PoCs for 1,055 (50.5%) vulnerabilities of the total dataset. Considering only honeypot data, we could find public PoCs for 317 out of 759 (58%) exploited vulnerabilities.

Overall, this shows that most vulnerabilities are exploited few times by few threat actors for a limited period of time and following exactly or very closely publicly available exploits. Only some vulnerabilities stand out and get adopted into botnets or automated scanners to be used more frequently and with more variation.

3. KEV remediation

KEY FINDING	WHY IT MATTERS	REMEDIATION
15% of exploited vulnerabilities have no patches available or have patches only available for some versions.	When a vulnerability cannot be patched, risk mitigation becomes more complicated, so the only option may be to disconnect a device.	Use segmentation and zoning to minimize network exposure to unpatched assets.

Although CISA mentions the existence of “a clear remediation action” as a criterion for inclusion in KEV that does not always mean that a vulnerability can be patched. In some cases, the product is discontinued and there are no patches available. In others, only some versions can be patched while other versions are considered end-of-life. In the cases where there is no patch, CISA recommends users to disconnect affected devices.

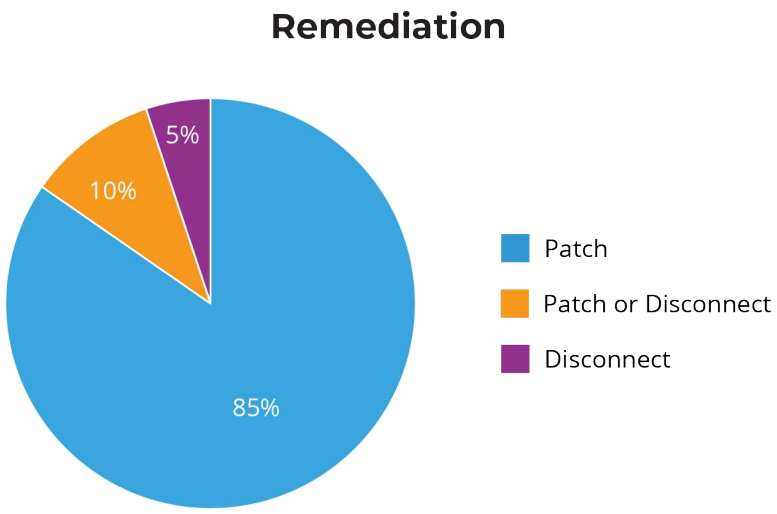

Figure 13 – Distribution of recommended remediation actions for vulnerabilities in CISA KEV

Figure 13 shows the distribution of recommended remediation actions for vulnerabilities in CISA KEV. The majority (85%) have available patches, but 15% either have no patches available or have patches available only for some versions.

3. Tips for defenders

Include other sources of vulnerability information than CVEs in the NVD for more complete risk assessments and to know what assets in the network may be currently vulnerable. Use multiple sources of information about exploited vulnerabilities to cover a larger threat landscape and prioritize the vulnerabilities that need to be patched.

CVSS scores do not always imply actual risk, but they are an important predictor of exploitation, so focus on mitigating vulnerabilities with higher scores, but it may leave some blind spots. Add extra layers of protection for web applications, such as web application firewalls. As with CVSS scores, vulnerability age does not always imply actual risk. Focus on mitigating more recent vulnerabilities but this may also leave some blind spots.

Prioritize assets with internet exposure when patching. For assets that cannot be patched, use segmentation to minimize their network exposure and likelihood of compromise without impacting mission-critical functions or business operations. Segmentation and zoning also limit the blast radius and business impact if a vulnerable asset becomes compromised. Use network and host detection systems and ensure that their signatures are updated. Focus on network detection for unmanaged devices such as OT/IoT:

- Monitor all network traffic for malformed packets that try to exploit known vulnerabilities or possible zero days.
- Anomalous and malformed IP traffic should be blocked, or at a minimum, network operators should be alerted to its presence.

4. Conclusion

Given the growing number of vulnerabilities found and exploited by malicious actors, organizations need help to understand what to prioritize. CISA KEV is an important resource to help with this prioritization by identifying vulnerabilities that have been or are being exploited, but it suffers from issues. They include a lack of transparency on the selection of vulnerabilities.

In this report, we showed that there are many exploited vulnerabilities not captured in CISA KEV that affect real organizations. In addition, we showed that no single database includes every exploited vulnerability, so organizations should rely on multiple sources.

Keep in mind that a list of exploited vulnerabilities is not useful if it cannot lead to risk mitigation. For example, a device running a vulnerable HTTP server version, but configured with that service disabled do not present an immediate risk. To implement risk mitigation in a timely and efficient manner, organizations need a way to:

- Automatically identify assets on a vulnerable network
- Identify issues currently being exploited
- Automatically understand the context on which these assets may be vulnerable

Only with the full picture can a database of exploited vulnerabilities be effectively used for patching prioritization and risk mitigation. Even CISA recommends that their KEV is used as an input to a process, such as the Stakeholder-Specific Vulnerability Categorization (SSVC). This is especially true for OT networks where patching is a hard and time-consuming effort that needs to be carefully planned.