# Time to next exploit

Stephen Robinson, Senior Threat Intelligence Analyst, WithSecure

# Time to next exploit

## Key Findings

Between January and May 2025:

- A vulnerability that went on to be exploited was published every 2 days, a 23% higher rate than in 2024.
- A zero-day vulnerability that was then found to be exploited was published every 3 days, a 46% higher rate than in 2024. This is double the growth rate of exploited vulnerabilities as a whole, suggesting either increasing targeting of zero-days by actors, or a failure to identify vulnerabilities by developers.
- A security service vulnerability was published every 10 days, a 34% rate of increase – This means that in 2025 security service vulnerabilities are being discovered at a 50% faster rate than average.
- A security service zero-day that was then found to be exploited was published every 15 days. While this is a 15% increase on 2024, that means the growth rate was lower than for zero-days as a whole, or even for all vulnerabilities. This suggests that security improvements in security service development may well be having an effect.

## Introduction

Organizations are facing an attack surface that is not only expanding at an unprecedented rate but also becoming more difficult to manage using traditional security approaches. The first half of 2025 has shown a sharp rise in both the discovery and exploitation of vulnerabilities, especially zero-days and those affecting security services, indicating that attackers are moving faster than defenders can respond. A new exploited vulnerability is published every two days, and a new exploited zero-day every three, with both categories growing significantly faster than in 2024. This growing gap between discovery and mitigation underscores a fundamental reality: reactive defence is no longer sufficient. Companies must adopt a proactive approach centred around continuous exposure management: monitoring, prioritizing, and remediating vulnerabilities before they are exploited. This research draws on verified exploitation data and real-world vulnerability trends to demonstrate why exposure management is no longer optional, but a foundational requirement for any organization aiming to stay ahead of cyber threats.

## Vulnerability data

Making accurate decisions requires accurate data, or at least as accurate as we can make it. In the world of vulnerabilities our best source of data is the CVE program and the NVD, at least for as long as they continue to be funded. The CVE program, and the CVSS data associated with CVEs, gives us a useful starting point to assess CVEs relative to each other, but we also need more real-world data. This is where the KEV comes in, as by definition it provides us with a list of vulnerabilities which we know have been exploited. The KEV is sometimes criticised for not being comprehensive enough and not including vulnerabilities which are reported to have been exploited. While we cannot comment on individual cases, the KEV exists to provide concrete data about exploitation and known threat, as such it most likely does have a restricted set of trusted inputs. This does at least mean that if a vulnerability is added to the KEV, we can have high confidence that it really has been exploited. Finally, there is the case of zero-day vulnerabilities. Whether a vulnerability that is a zero-day is not tracked by the KEV, the NVD, or the CVE program. This is not to say they do not include zero-days once they are known about, but they do not track whether a vulnerability was or was not a zero-day. Fortunately, that data is tracked elsewhere, such as at zero-day.cz which was referenced for this research.

One interesting thing to consider is that the developers of a piece of software could be considered to have more control over whether a zero-day is exploited than they have over whether an n-day (i.e. a non-zero-day) is exploited. This is because the only way to prevent a zero-day being exploited is to not introduce it into your software, which is entirely on the developers. There is no possibility of patching, the responsibility is entirely on the developers to be preventative. However, with an n-day patching is possible, yet while developers can make a patch available they cannot in the most part force users to apply it. This then implies that part of the responsibility for exploitation of n-days rests on those who for whatever reason did not apply the available patch, which does seem like a valid conclusion.

# Method

The numbers used here only include CVEs published in 2023, 2024, and 2025. Not vulnerabilities from previous years, and it uses **the date the CVE was published, not the date it was added to the KEV**. This removes the large numbers of n-days from previous years which could be exploited and added to the KEV  and makes the data more relevant to the current time frame. The reason to exclude CVEs published in earlier years is that their existence speaks to the state of historic security, not necessarily the current state of security. By only looking at vulnerabilities by the date they were published, we can get a better idea of current software security, and make closer comparisons between n-days and zero-days, based on the assumption that zero-days are generally only going to be exploited for a short period of time before they are identified, as opposed to going undetected while being exploited for multiple years.

By combining KEV, NVD, and zero-day data sources, we can then draw out statistics and perform comparative analysis to inform our decision making around the current state of vulnerabilities and exposure. This is something we did in early 2024, when we found that exploitation of vulnerabilities affecting edge services and infrastructure devices was on the rise.  So now, let’s look at the data for 2024/25 and see what may have changed, and what new insights we can derive.

# KEV vulnerabilities

From the beginning of January 2025 until the end of May, there were 15 vulnerabilities published and per month that were added to the KEV, or roughly 1 vulnerability every 2 days. That’s an increase of 23% on 2024:

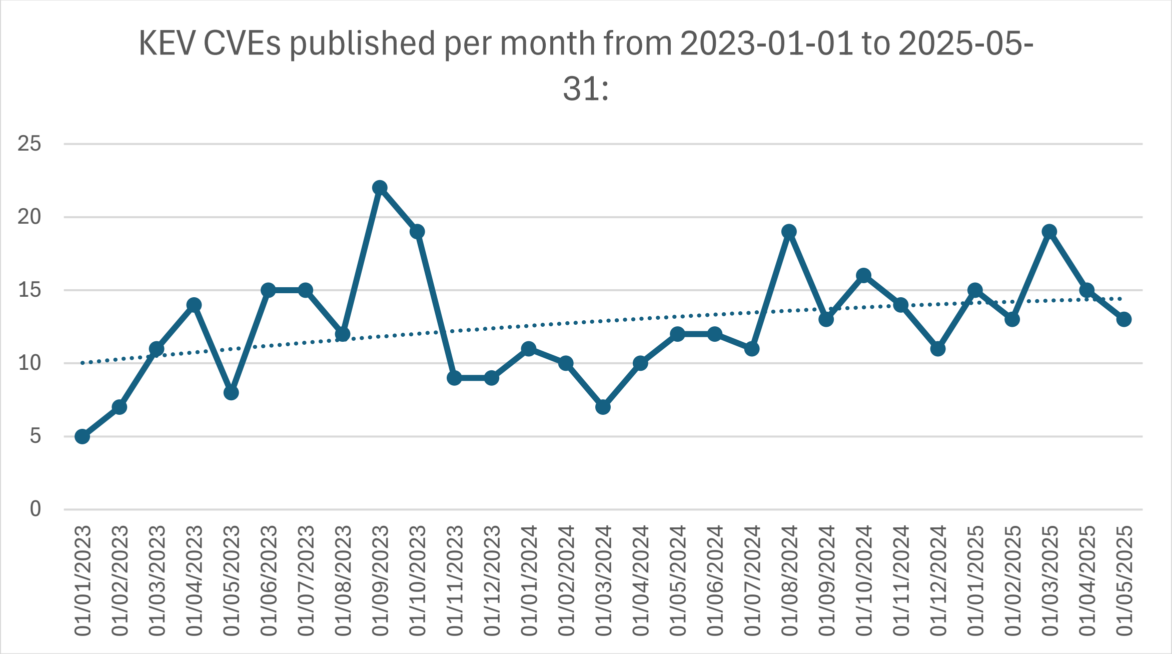

The severity of published vulnerabilities has not changed significantly over time, though this may be explained by the fact that the result of the CVE severity calculations is as much to do with the vulnerable software and how it can be exploited, rather than the individual exploit itself. For example, however severe an exploit, you cannot exploit the network interfaces of a product if it simply does not have any.

## Zero-days added to KEV

Unfortunately, if we look at zero-day vulnerabilities we find that so far in 2025, 9.6 zero-day vulnerabilities found to have been exploited have been published each month, a quite shocking rate of one every 3 days, or just over 2 per week:

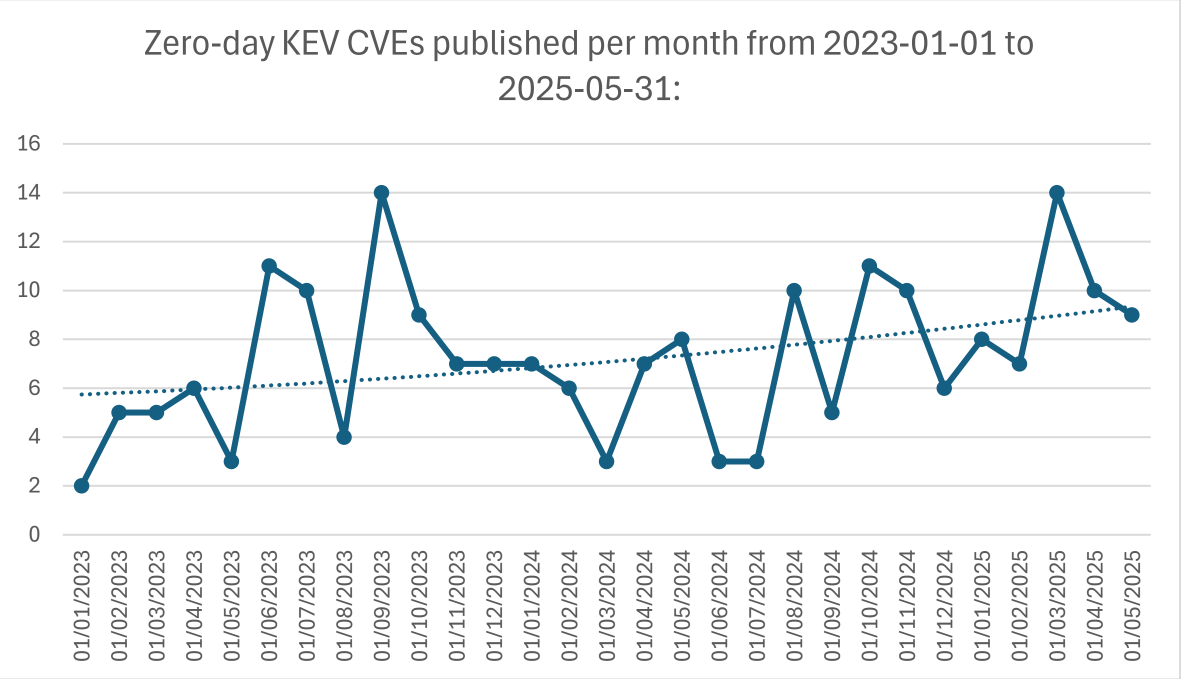

That is a 46% increase from 2024, meaning that the rate of increase of known exploited zero-days in 2025 is twice the overall rate of increase for known exploited vulnerabilities:

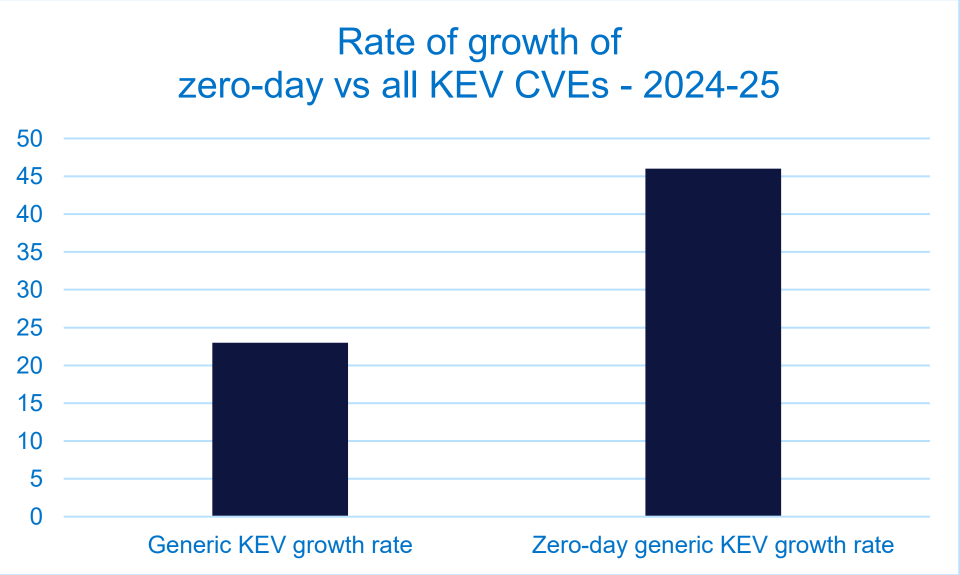

It is interesting that the rate of growth of zero-days is higher than that of non-zero-days, or n-days. A zero-day is a vulnerability which is known and can be exploited by attackers, but has no available patch, while an n-day is a vulnerability for which a patch is available before an exploit. As such, the difference really comes down to who finds it and addresses it first – Attackers or defenders/developers.

As such, higher zero-day than n-day growth rates are likely down to a combination of two things:

- Increased attacker focus on discovering and exploiting zero-days.

- A lack of investment in code quality checks and security assessment of software products by developers.

## Security service vulnerabilities

Of course, there are ways to defend against zero-days, such as implementing a defence in depth strategy (now more fashionably referred to as zero-trust) of multiple security tools. But how secure, on average, are those security tools?

There were 3.4 security service-related CVEs (Anti-virus, VPN gateway, Firewall, etc.) published per month and added to the KEV in 2025, which is a 32% increase on 2024.

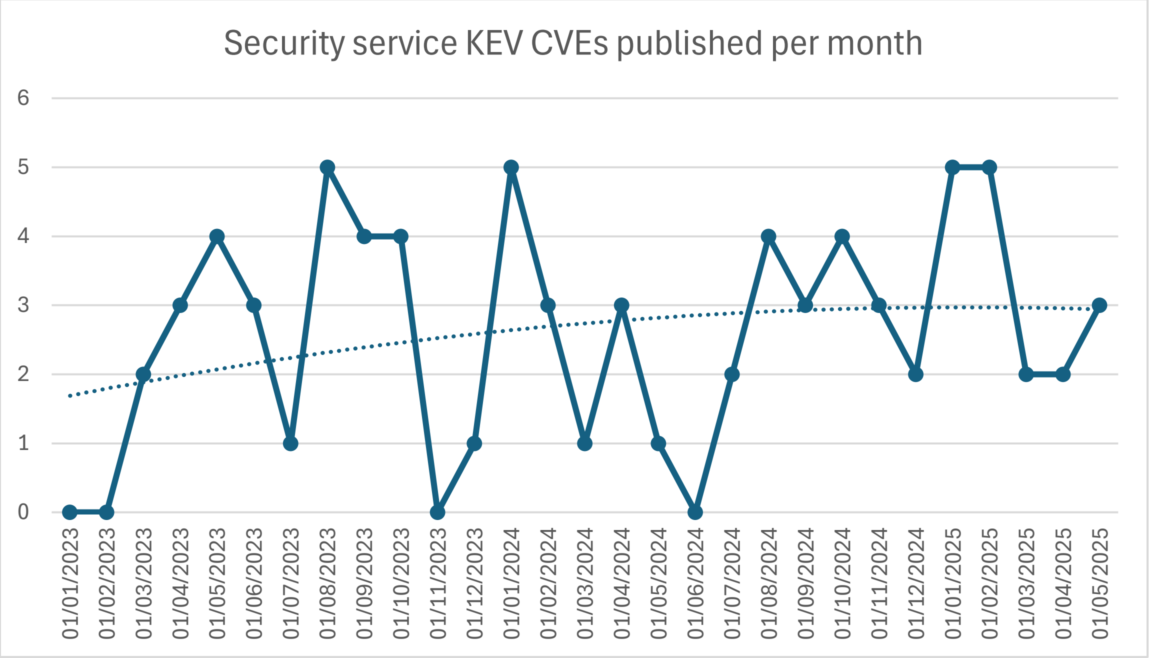

That gives us a rate of increase that is 50% higher than KEV CVEs as a whole:

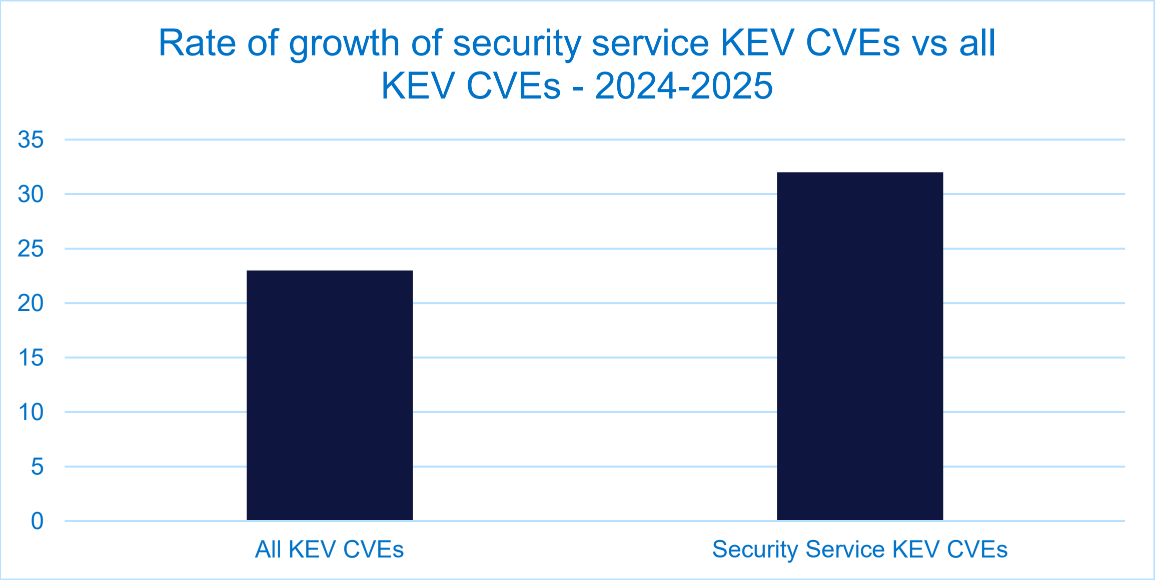

A rate of growth which is higher than the average is a bad thing and suggests either increased targeting of this type of software, or that the software being released is itself more vulnerable than the average.

However, if we look at the number of zero-days in security service software for 2025, while this did increase to a rate of 2.2 zero-days per month, that is only a rise of 15% on 2024. As such, the rate of growth of security-service zero-days is lower than that of zero-days as a whole, or of security service vulnerabilities as a whole:

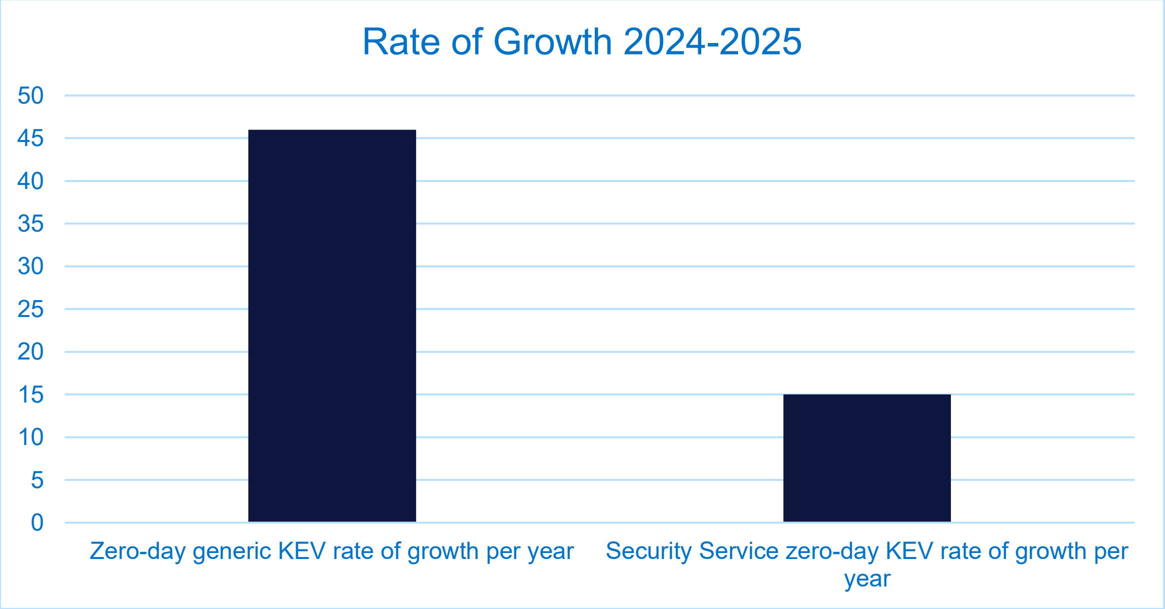

If we consider that the responsibility for zero-days lies solely on the developers of a piece of software, the lower rate of growth of exploitation of zero-days compared to exploitation of n-days might suggest that the software development practices of security software developers are improving relative to the average, but that patching processes for n-days are not, thus increasing the chance that n-days will be exploited, while decreasing the chance that zero-days will exist and be exploited.

# Edge services

The rate of growth of known exploited edge service CVEs from 2024 to 2025 was 13%, increasing from 5.33 to 6 per month. Meanwhile, known exploited edge service zero-days in that time increased 20% from 2.17 to 2.6 per month. That is around a 150% higher rate of growth than the non-zero days, though in both cases these are lower than the overall growth rates for all vulnerabilities and zero-days:

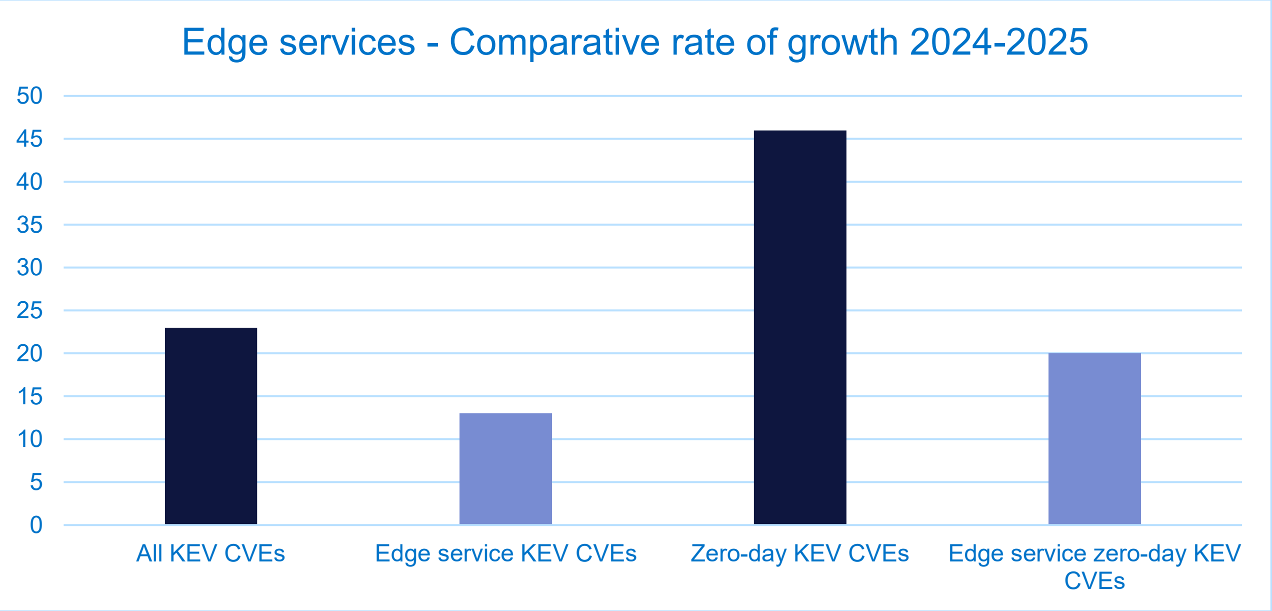

Known exploited infrastructure vulnerabilities actually show an even lower rate of growth than edge services, with a 9% increase from 3.5 per month in 2024 to 3.8 per month in 2025, much less than half the average rate of growth for generic KEV vulnerabilities. Known exploited zero-day infrastructure vulnerabilities show a growth rate of 20% this year, similar to zero-day edge service vulnerabilities. These similar growth rates between edge and infrastructure type vulnerabilities is likely because the two categories have a tendency to overlap - many network edge devices are also infrastructure.

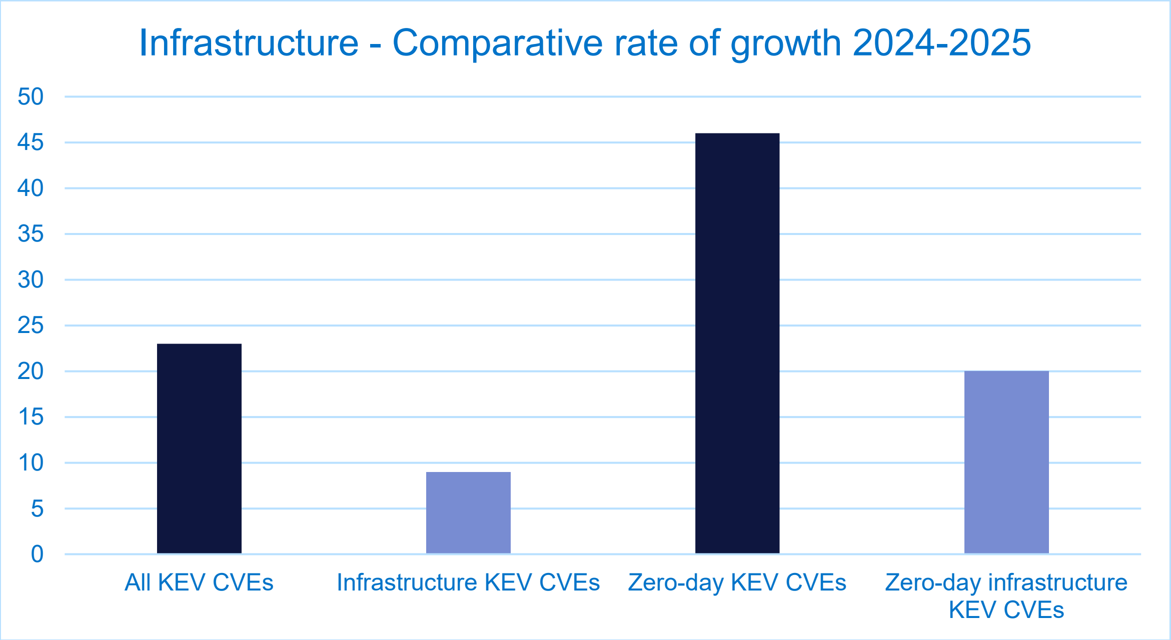

The 20% increase in infrastructure zero-days from 2024-2025 is a slower rate of growth than was seen in 2023-2024, when the number of zero-days per month increased by 50%, from 1.33 to 2:

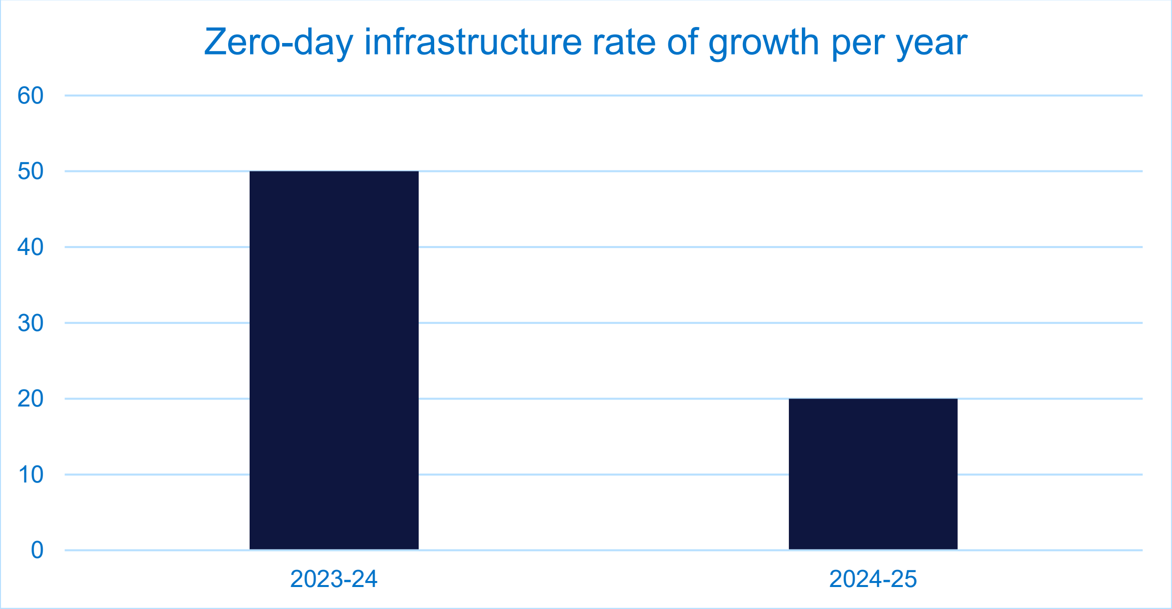

Now, it's good to see the rate of growth going down, but the number of vulnerabilities being exploited is still increasing each year:

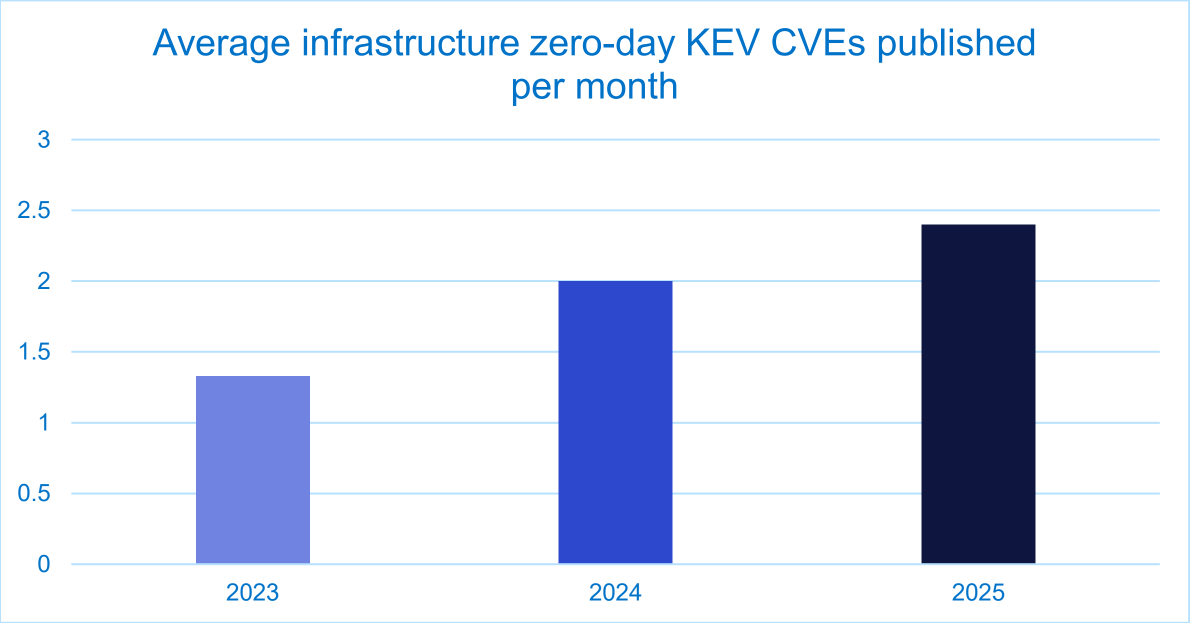

There is no reason to believe that attackers are targeting these vulnerabilities less, or have less appetite for zero-days, which suggests that the increased efforts of developers and the wider security industry are having a positive effect, however, those efforts need to continue and increase if we wish to actually become more secure.

## 2023 – 2024

When dealing with computers, we are quite used to numbers going up every year, and unfortunately that is often the case when we are dealing with cyber-crime as well. Reading the statistics here we could be forgiven for thinking that the number of exploited vulnerabilities and zero-days increases year on year, and it is simply a question of how much. However, that is not necessarily the case. In 2023 and 2024 there was the exact same number of CVEs published in those years added to the KEV, 146, or 12.17 per month:

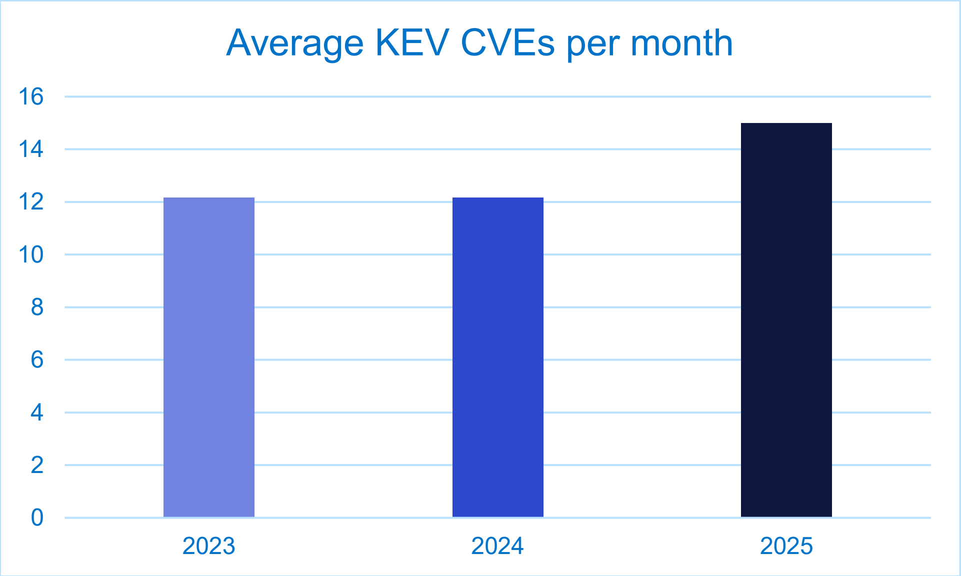

Looking at zero-days, there was actually a 5% drop in the number of zero-days added to the KEV in 2024 compared to 2023:

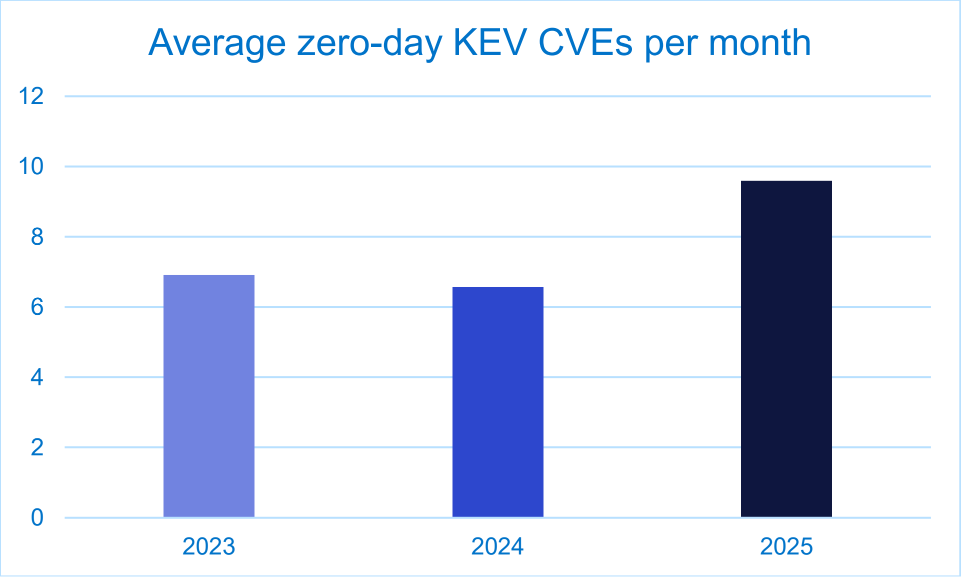

This really shows up the 46% growth rate for zero-days added to the KEV from 2024 to 2025, especially as there is no indication of any kind of annual Q1 surge in volume, meaning that this large increase is unlikely to have been skewed by looking at only the first 5 months of the year:

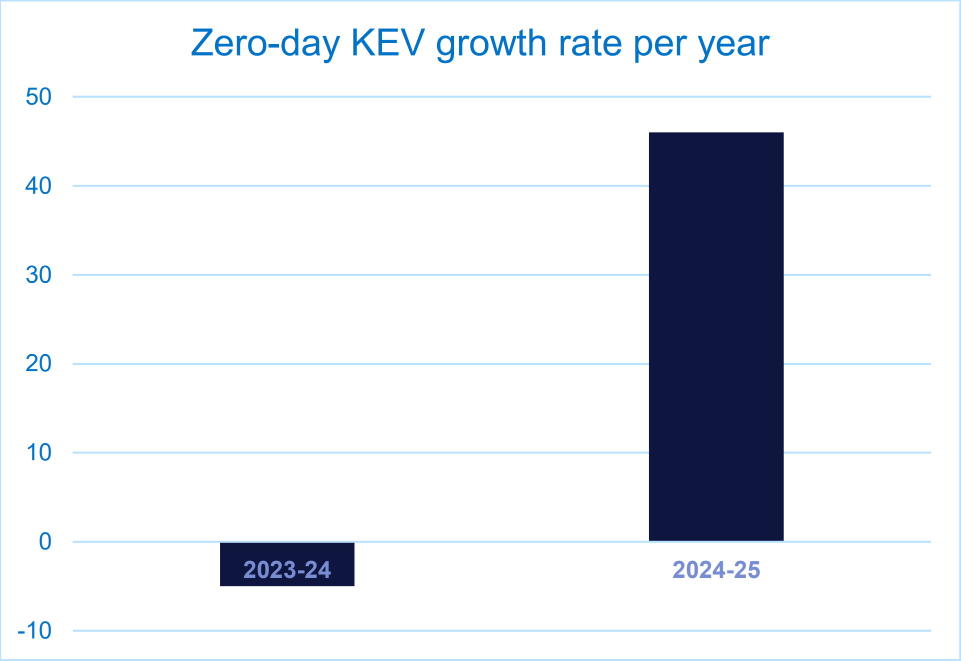

## Conclusion

Now, I love a good statistic, and I love exploring the threat landscape that we as defenders face, but unfortunately, this brief yet vicious whirlwind of statistics on its own does not actually tell you how you should be defending your organization. We all know that we should be applying patches and following security best practice already, but we also know it’s not that simple. Teams have technical debt, networks have unexplored, dark corners, some devices cannot be patched without major downtime, and in some cases patches simply are not available. In the real world, work must be prioritized, and prioritization requires information and awareness.

This is where exposure management comes in. Exposure management gives an organization visibility of its attack surface in the form of information about the most urgent and impactful vulnerabilities, weaknesses, or misconfigurations that need to be prioritized. It can also facilitate and automate that work – after all, there’s nothing worse than being told about all the work that your team hasn’t yet done, while knowing that you don’t have time to do it. Good exposure management ensures that informed strategic and tactical decisions can be made, and reduces the time and effort needed to secure a network by providing the automation needed to enable 24/7, machine speed responses. Having exposure management doing the hard work for you is far better than trying to spot and address a vulnerability every 2 days.