BeyondTrust

# 2026 Microsoft Vulnerabilities Report

## 13th EDITION

Data-packed insights and expert analysis to help you mitigate security risks in your Microsoft estate.

BeyondTrust

## TABLE OF CONTENTS

EXECUTIVE SUMMARY 2026 - 13TH EDITION

# Uncovering the Ghost in the Machine

In this 13th edition of the BeyondTrust Microsoft Vulnerabilities Report, the threat landscape is anything but ordinary.

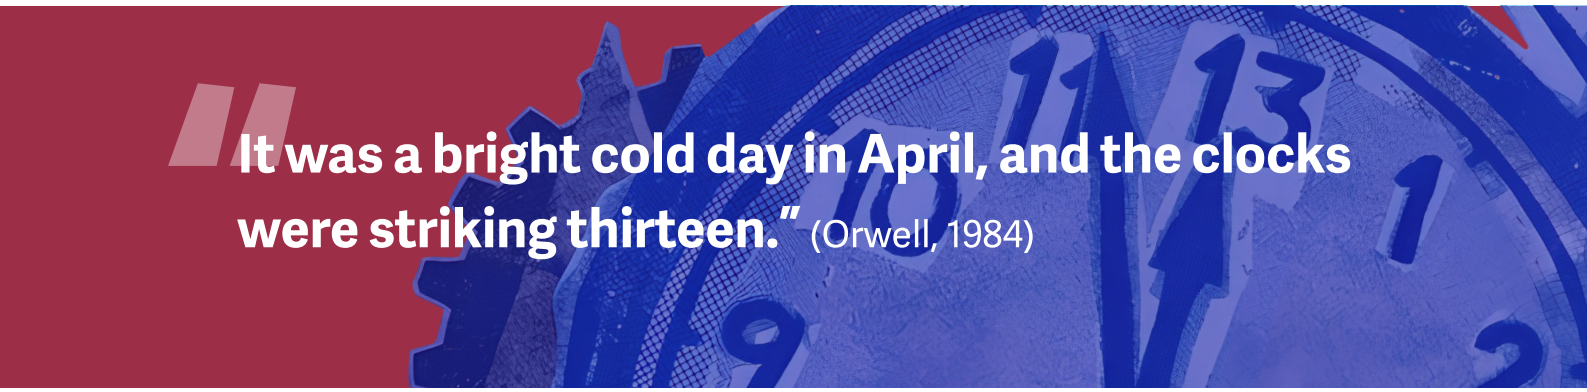

> **"It was a bright cold day in April, and the clocks were striking thirteen."** (Orwell, 1984)

A great technology shift is underway. In his novel 1984, George Orwell describes a "Versificator"—a mechanical device used to produce literature and music without human intervention. Now, in this age of AI, machines can churn out that "dreadful rubbish" Orwell described. Machine content that's entertaining and enraging in equal measure is no longer fiction, but a daily reality.

Aside from the Orwellian clock reference, thirteen is an infamous number, associated with bad luck, unexpected twists, and stories that take a darker turn when you least expect it. In this report, we're leaning into that theme because 2025's data has exactly that shape: familiar on the surface, but disquieting the deeper you look.

At first glance, we see the total number of Microsoft vulnerabilities declined slightly year-over-year, from 1,360 in 2024 to 1,273 in 2025, continuing a minor fluctuation pattern in place since 2020. While the overall numbers remain significantly higher than the 540 we saw 10 years ago, that disparity is largely explainable by the expansion of the Microsoft software portfolio. We must also consider that Microsoft has increased focus on security in recent years, and, as a result, more vulnerabilities have been found and added to the overall count.

But once we dissect the data, things take an unexpected turn. Yes, total volume of vulnerabilities dipped by about 6%, but critical vulnerabilities doubled, **rising from 78 to 157.** This represents a considerable departure from the steady, multi-year decline in the volume of critical vulnerabilities.

This matters because defenders don't lose sleep over raw vulnerability counts; they lose sleep over what those critical vulnerabilities enable: privilege escalation, lateral movement, identity abuse, and rapid compromise across connected systems. And the 2025 data reinforces the same hard truth: Elevation of Privilege (EoP) remains the dominant vulnerability category, accounting for 40% (509) of all vulnerabilities.

That's not a coincidence. It's a signal that the biggest "amplifier" in many real-world incidents is still the same: once an attacker gets a foothold, the fastest path to business impact is often to quietly escalate privilege and operate as a trusted identity—which, in 2026, could just as easily be an AI agent or machine identity as a human one.

In modern identity attacks, exploiting software vulnerabilities is just one of many ways to elevate privilege, alongside excessive standing privilege, identity infrastructure misconfigurations, and weak identity controls.

For IT practitioners, remediating vulnerabilities, such as by patching systems in a timely manner, reduces the window of opportunity for exploitation, thus lessening the likelihood of a breach. However, patching must be implemented alongside the principle of least privilege (PoLP). This is a key preventative capability for maintaining a hardened security posture, allowing organizations to withstand attacks, while mitigating both software and identity risks. A least privilege approach also goes a long way toward minimizing the "blast radius" in the event of a zero-day exploitation or a compromised identity.

In addition to our usual data slicing and dicing, this year's report also shines a light on why many cloud and AI vulnerabilities never become CVEs, what the explosion of agentic AI means for our attack surface, and what's going on with the modern threat landscape.

As in years past, we've assembled a panel of some of the world's leading cybersecurity experts to weigh in on the report findings. These cyber heroes will help us collectively set our sights forward on emerging threats, new vulnerabilities, and best practices for building cyber resilience across the enterprise, and society at large.

Read on to better understand, identify, and address the risks facing your organization within the Microsoft ecosystem.

# Key Findings and Data

## Key Findings

Total vulnerabilities dropped slightly YoY, from 1,360 to 1,273. This follows a trend of minor fluctuations since 2020. However, the number of critical vulnerabilities doubled last year, rising from 78 (2024) to 157 (2025).

Total vulnerabilities in Azure and Dynamics plateaued in 2025, at 69, while critical vulnerabilities rose sharply from 4 to 37—hitting a record high. Windows Server total vulnerabilities increased 14%, up from 684 to 780, while critical vulnerabilities increased by 16%.

Microsoft Office total vulnerabilities surged 234% year-on-year, from 47 to 157, and critical vulnerabilities rocketed from just 3 to 31 over the same period. On the other hand, total vulnerabilities of Microsoft Edge reached an all-time low, diminishing substantially from 292 to just 50 (83% decrease).

In terms of vulnerability types, the Elevation of Privilege category continued to dominate in 2025, accounting for 40% of all vulnerabilities. After rising in 2024, both Remote Code Execution and Security Feature Bypass saw declines in 2025, while Information Disclosure vulnerabilities increased by 73% last year.

## Data Highlights

- Total vulnerabilities in 2025 decreased by 6% to **1,273** (down from 1,360 in 2024)
- The number of **critical vulnerabilities doubled year-over-year,** from 78 to 157
- Continuing a post-pandemic trend, the Elevation of Privilege category accounted for a massive **40% (509) of the total vulnerabilities** last year
- Microsoft Azure and Dynamics 365 **experienced a 9x rise in critical vulnerabilities, from 4 to 37**
- Microsoft Edge experienced 50 vulnerabilities last year, **83% less YoY**
- There were **612 Windows vulnerabilities published in 2025,** 36 were critical
- Windows Server had **780 vulnerabilities in 2025, 50 were critical**
- Microsoft Office experienced **157 vulnerabilities** in 2025, over 3x the total recorded in 2024. Critical vulnerabilities experienced a **10x increase from last year.**

# Understanding Microsoft's Classification of Critical Vulnerabilities

## How Does Microsoft Classify Critical Vulnerabilities?

Microsoft significantly updated how they classify critical vulnerabilities in 2020 by leveraging the National Vulnerabilities Database (NVD) Common Vulnerability Scoring System (CVSS), which classifies critical vulnerabilities as those awarded a score of 9.0 - 10.0. This brought them into line with industry standards. However, as the system lacks nuance around the real-world impacts of a given vulnerability, Microsoft also uses their own Security Update Severity Rating System, which rates each vulnerability according to the worst theoretical outcome, should that vulnerability be exploited. In March 2025, they updated from CVSS 3.1 to CVSS v4 as the preferred scoring standard for vulnerabilities. With both CVSS v4 and Microsoft's severity rating system, users can see a more multifaceted picture of what defines a 'critical' vulnerability.

### This means that:

This year, 42 Microsoft vulnerabilities were scored as "critical" under the National Vulnerability Database (NVD) scoring system (about an 8% increase from 39 in 2024). Yet, under Microsoft's Security Update Severity Rating System, "critical" vulnerabilities increased by 101%, from 78 to 157.

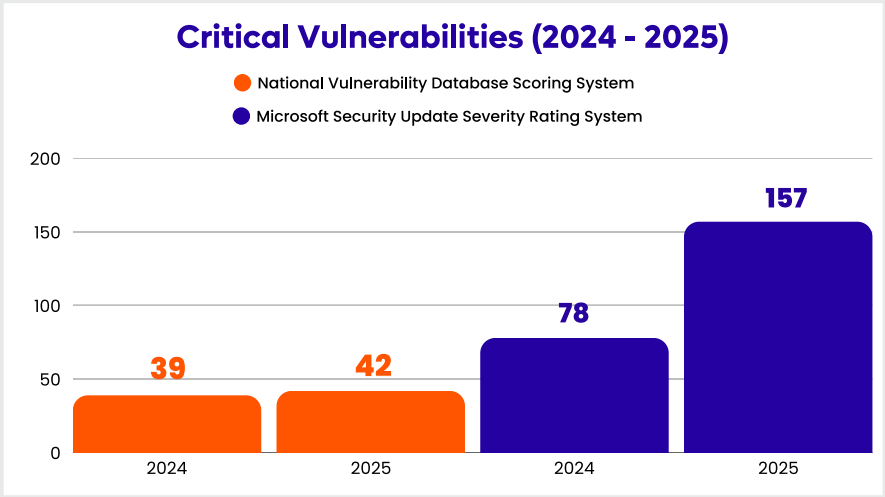


Figure 1: Microsoft saw an increase to 42 "critical" vulnerabilities in 2025, as classified based on their CVSS scores, while Microsoft's Security Update Severity Rating System reported a significantly higher number of "critical" vulnerabilities, at 157.

**CVSS v4 Ratings**

| Severity | Base Score Range |
|---|---|
| None | 0.0 |
| Low | 0.1 - 3.9 |
| Medium | 4.0 - 6.9 |
| High | 7.0 - 8.9 |
| Critical | 9.0 - 10.0 |

Figure 2: The National Vulnerabilities Database (NVD) scoring system for classifying critical vulnerabilities

**Microsoft Security Update Rating System | Rating & Description**

**Critical**

A vulnerability whose exploitation could allow code execution without user interaction. These scenarios include self-propagating malware (e.g. network worms), or avoidable common use scenarios where code execution occurs without warnings or prompts. This could mean browsing a web page or opening an email. Microsoft recommends that customers apply critical updates immediately.

**Important**

A vulnerability whose exploitation could result in compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources. These scenarios include common use scenarios where a client is compromised with warnings or prompts, regardless of the prompt's provenance, quality, or usability. Sequences of user actions that do not generate prompts or warnings are also covered. Microsoft recommends that customers apply important updates at the earliest opportunity.

**Moderate**

Impact of the vulnerability is mitigated to a significant degree by factors such as authentication requirements or applicability only to non-default configurations. Microsoft recommends that customers consider applying the security update.

**Low**

Impact of the vulnerability is comprehensively mitigated by the characteristics of the affected component. Microsoft recommends that customers evaluate whether to apply the security update to the affected systems.

Figure 3: The Microsoft Security Update Severity Rating System for identifying associated risk based on the worst theoretical outcome, if that vulnerability were to be exploited

The difference between CVSS scoring and Microsoft's severity rating system is worth noting—not only when considering the data in this report, but also when considering risk in your organization. CVSS scores measure the technical severity of a vulnerability (i.e., whether a vulnerability results in some loss of data confidentiality or total loss of confidentiality). CVSS scoring does not measure the risk of that vulnerability.

Microsoft's severity rating system is potentially far more useful than a base or temporal CVSS score for security practitioners trying to prioritize risk reduction. However, it's important to know your own environment and risks so you can understand how best to prioritize patches and / or use other security hardening controls and mitigations.

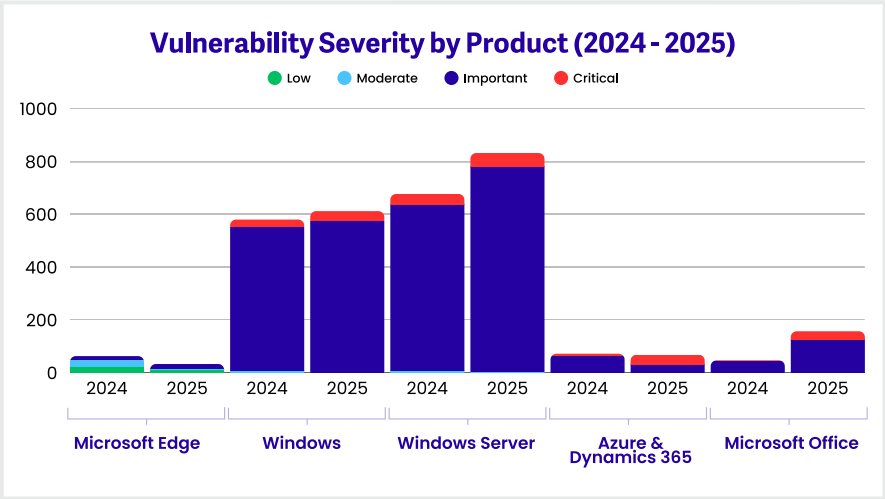


Figure 4: Microsoft Security Update Severity Rating System's breakdown of vulnerability severity by product shows patterns that are not visible when only looking at total vulnerability count. We see a significant increase in 'important' impact vulnerabilities in Windows Server, and in both 'important' and 'critical' impact vulnerabilities in Microsoft Office. Azure and Dynamics 365 have a similar number of total vulnerabilities when comparing 2024 to 2025, but show a notable increase in 'critical' impact vulnerabilities.

All the data provided by Microsoft's severity rating system is based on the information available at the time. It lacks the context of your own organization's threat models. What is considered a critical patch for one organization may not be so for another organization; it all depends on the business context.

In addition to their Security Update Rating System, Microsoft also publishes an Exploitability Index to help customers understand the likelihood of exploitation and prioritize the most pressing security updates for their Microsoft environments. As a word of caution, this information reflects the likelihood of exploitation at the time the security update was published. It may not reflect the real-world exploitability that develops in the following weeks or months as more threat actors weaponize the vulnerability. The index is best used for very short-term prioritization of updates, rather than as a justification to significantly delay patching.

**Microsoft Exploitability Index**

| Index | Short Definition | Expanded Definition |
|---|---|---|
| 0 | Exploitation Detected | Microsoft is aware of an instance of this vulnerability being exploited. As such, customers who have reviewed the security update and determined its applicability within their environment should treat this with the highest priority. |
| 1 | Exploitation More Likely | Microsoft analysis has shown that exploit code could be created in such a way that an attacker could consistently exploit this vulnerability. Moreover, Microsoft is aware of past instances of this type of vulnerability being exploited. This would make it an attractive target for attackers, and therefore, more likely that exploits could be created. As such, customers who have reviewed the security update and determined its applicability within their environment should treat this with a higher priority. |

Figure 5: The Microsoft Exploitability Index for assessing the likelihood of exploitation at the time the security update was published

| | | |
|---|---|---|
| 2 | **Exploitation Less Likely** | Microsoft analysis has shown that, while exploit code could be created, an attacker would likely have difficulty creating the code, requiring expertise and / or sophisticated timing, and / or varied results when targeting the affected product. Moreover, Microsoft has not recently observed a trend of this type of vulnerability being actively exploited in the wild. This makes it a less attractive target for attackers. That said, customers who reviewed the security update and determined its applicability within their environment should still treat this as a material update. If they are prioritizing against other highly exploitable vulnerabilities, they could rank this lower in their deployment priority. |
| 3 | **Exploitation Unlikely** | Microsoft analysis shows that successfully functioning exploit code is unlikely to be utilized in real attacks. This means that, while it might be possible for exploit code to be released that could trigger the vulnerability and cause abnormal behavior, the full impact of exploitation will be more limited. Moreover, Microsoft has not observed instances of this type of vulnerability being actively exploited in the past. Thus, the actual risk of being exploited from the vulnerability is significantly lower. Therefore, customers who have reviewed the security update to determine its applicability within their environment could prioritize this update below other vulnerabilities within a release. |

Overall, these various metrics and measurements remind us that there are many different contextual ways to break down data and understand risk. Raw CVSS only makes up one piece of the puzzle; on its own, it cannot tell you whether a vulnerability is actively being exploited in the wild, if it could cause a mission-critical impact on a system, or if the limited loss of some highly sensitive data would have a more severe impact than the total loss of less-sensitive data.

This is one of the never-ending challenges of patch management; there isn't a one-size-fits-all approach, and it's as much about knowing your own organization as it is about understanding the threat landscape.

# 5-Year Trend

## When There's More Than Meets the Eye

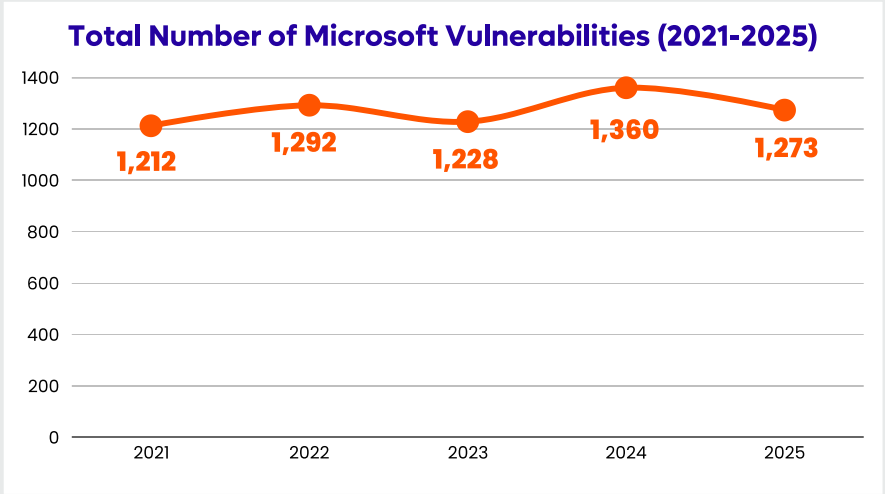


Figure 6: Total number of Microsoft vulnerabilities remained fairly stable—dropping by 6% compared to 2024.

Fellow data nerds may have noticed a hidden **mean**ing in this data: the 1,273 number for 2025 is coincidentally the same as the mean average of the previous 4 years. So, when it comes to the total number of vulnerabilities, we can certainly call it an average year. Overall, it's good to see the numbers drop by 6%, back into the 1,200 range, after an 11% increase the previous year.

While stability is welcomed, it shows that Microsoft's Secure Future Initiative (SFI), which launched in 2023 to transform their security culture and link executive pay to security outcomes in the light of modern threats, still has a ways to go.

Long-time industry veterans will have detected the echo of the 2002 secure computing memo, in what appears to be a 20-year cycle of refocusing on security. As the 10-year trend below illustrates, we've been on an upward roller coaster ride for some time, holding our breath in anticipation of the big drop.

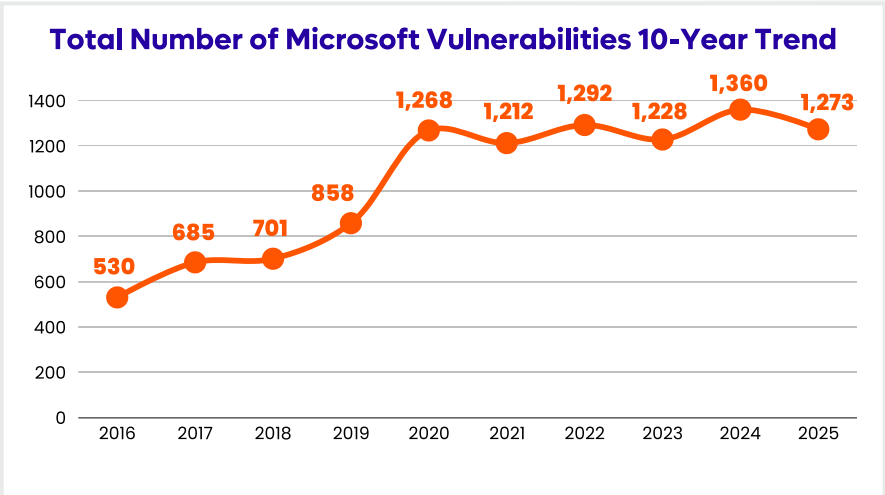


Figure 7: There is a pattern of relative stability when looking at a 10-year trend of total number of Microsoft vulnerabilities.

**We offer two potential interpretations of this relative stability:**

1. Given the rapid transformation that's occurring in this era of AI, the stability is a positive development and reassurance that Microsoft is doing a good job of securing their products, despite an expansion of the potential attack surface.

2. Given that many AI vulnerabilities and exploits don't get assigned a CVE or patch to deploy, and that Microsoft has been putting security front and center with the SFI, this number in total vulnerabilities should be decreasing, rather than remaining stable.

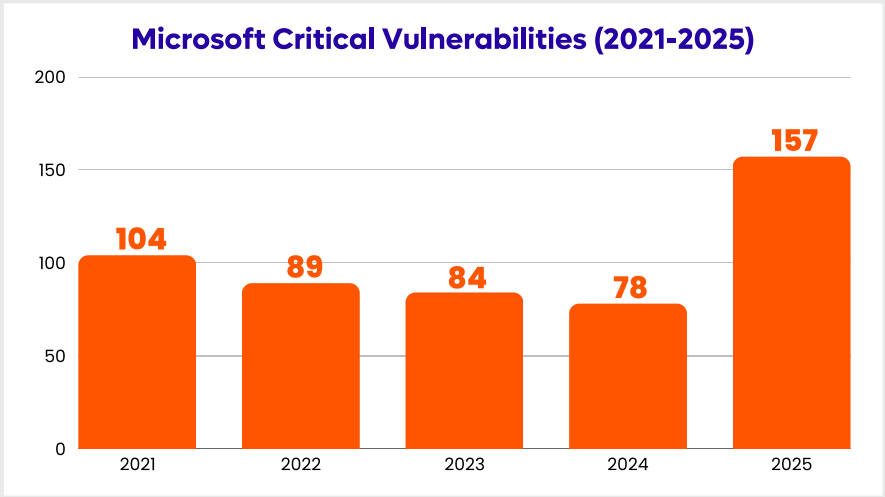


Figure 8: Critical vulnerabilities showed a significant increase in 2025, doubling from the previous year (78 to 157).

As our Chief Security Advisor, Morey Haber, says, "When talking to a doctor, you want to hear the words stable or improving—never getting worse. The same is true for CISOs".

It would be fair to call the total number of vulnerabilities 'stable', but when it comes to critical vulnerabilities, you may want to sit down for this news, as we are seeing a counter trend: a doubling of critical vulnerabilities from 2024 to 2025, as show in Figure 8 above.

With critical vulnerabilities representing around 12% of the total Microsoft vulnerabilities, all this is nowhere near the dark ages of 2013, where 44% of all Microsoft vulnerabilities were critical. However, it's a significant departure from long-term Microsoft trends we've observed for high-impact vulnerabilities.

Office, and Azure and Dynamics 365, are largely driving the critical vulnerability increases, while areas such as Edge are showing a decrease, and Windows / Windows Server remain relatively unchanged.

In the next section, we will explore how both critical and total vulnerability counts are showing up throughout different product lines, along with which categories of vulnerabilities we are consistently seeing.

# Vulnerabilities Data Deep Dive

## Vulnerabilities by Product Line

This year, we see a lot of vulnerability variances between Microsoft product lines. While some product areas appear to be taking positive steps, others are carrying the bulk of that significant critical vulnerability number mentioned earlier. Without further ado, let's dive into the vulnerability details for each product line.

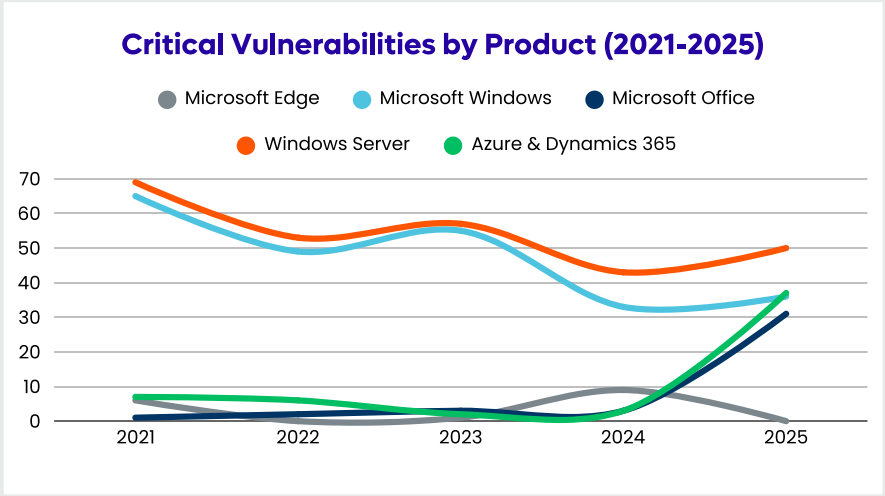


Figure 9: The bulk of Microsoft critical vulnerabilities are concentrated in a spike for Microsoft Azure & Dynamics 365 and Microsoft Office, alongside consistently high YoY numbers, within Windows Server and Microsoft Windows.

Windows Desktop and Server OS have historically been the dominant source of critical vulnerabilities. However, over the years, the vulnerability numbers for these product lines have diminished. While there is some year-on-year variability, the five-year trend clearly demonstrates a significant reduction. As we've discussed previously, this is partly due to the combination of older versions of Windows reaching end of support, and the security improvements offered in Windows 10 and 11.

Vulnerabilities for Microsoft Office, as well as for Azure and Dynamics 365, seemed on the verge of extinction last year. In 2025, this trend dramatically reversed course, but when it comes to vulnerabilities, as LL Cool J says, "Don't call it a comeback, I been here for years." A lack of vulnerabilities is sometimes the quiet before the storm. In previous editions of this report, we've discussed the snowball effect that happens when the discovery of one vulnerability causes more focus to be cast on an area that hadn't been put under the microscope in a while. This results in a flurry of activity and a spike in the data as more vulnerabilities are unearthed due to the scrutiny.

# Dark Clouds Gather

Time to get out the umbrellas. While the total number of vulnerabilities in Microsoft Azure and Dynamics 365 dipped slightly, critical vulnerabilities spiked dramatically, jumping from 4 to 37 in a single year. To put that in context, only 6% of the overall vulnerabilities were critical for this product line in last year's report. In 2025, critical vulnerabilities represented a 54% share of the total, serving as a sharp reminder to not just look at the totals.

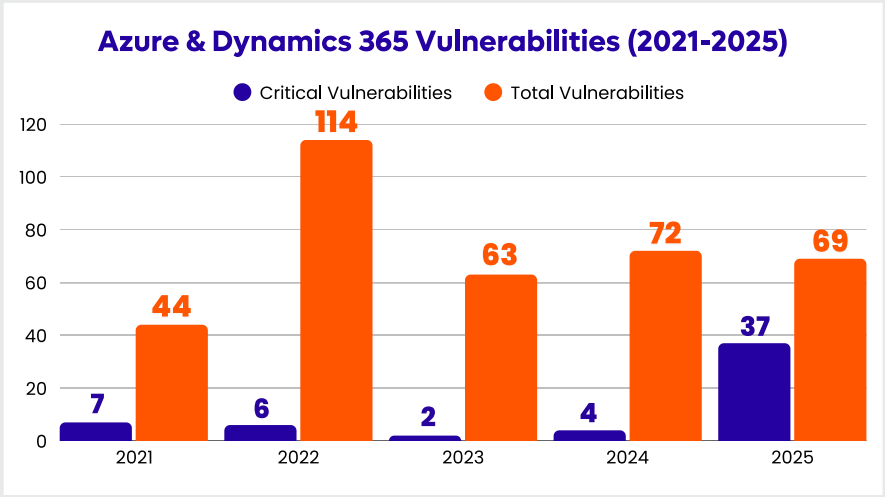


Figure 10: While total vulnerabilities in Azure and Dynamics 365 decreased slightly, critical vulnerabilities increased significantly, representing 54% of all vulnerabilities.

The concern for many is that cloud platforms are no longer just about infrastructure. They are crucial to business operations, providing a wide variety of services for identity and access management, business automation, etc., as well as serving as control planes for the entire enterprise's security. A critical flaw in these environments does not simply expose data; it can take down an entire workflow and, ultimately, business operations. It can collapse trust boundaries, and when cloud vulnerabilities turn critical, the blast radius becomes the defining risk metric.

### When you lose the lid for the container – CVE-2025-65037

Organizations are increasing their use of Azure Container Apps because it lets them run containerized apps while maintaining less infrastructure, with built-in scaling that can match demand and reduce cost. CVE-2025-65037 is a critical remote code execution vulnerability in those Azure Container Apps.

In the real world, these apps are often customer-facing services, background jobs, and automation that keep the business moving. This means a vulnerability allowing remote code execution isn't just another bug. It becomes a direct path into the workloads that process sensitive data, call internal services, and interact with the rest of the cloud environment. For these reasons, this particular CVE garnered the maximum CVSS score of 10.0.

CVE-2025-65037 is a prime example of the broader shift in risk as part of the cloud transformation story. It serves as a concrete illustration of why critical cloud vulnerabilities really matter, even when overall vulnerability counts look small in comparison to Windows.

When a cloud vulnerability is critical, it can wreak disproportionate impact, because cloud platforms act like control planes for modern business operations, and workloads often rely on highly privileged, non-human identities (NHIs) that authenticate and continuously call APIs. Exploitation and impact can occur at cloud scale and machine speed, increasing pressure to not only secure the cloud systems, but also control the blast radius of the NHIs driving them.

**One token to own them all - CVE-2025-55241**

CVE-2025-55241 is one of those rare vulnerabilities that makes you stop and think, "did I read that right?"

As a critical vulnerability in Azure Entra ID, CVE-2025-55241 could have enabled an attacker to impersonate another user, including a Global Administrator, in any organization's tenants. Due to the severity, Microsoft gave it the maximum CVSS score of 10.0 and quickly patched it on July 17, 2025, with no indication of exploitation in the wild.

In many companies, Entra ID is the front door for Microsoft 365 and Azure. If someone can impersonate any identity, especially a highly privileged one, they are no longer picking a lock on a single app; they are walking through the building with a master key. The researcher who uncovered this, Dirk-jan Mollema, described it as a breakdown in tenant separation, caused by a combination of internal actor tokens and a weakness in the Azure AD Graph API. In short, you could create an access token in one tenant that would be accepted in any other tenant in a way it never should have been.

This vulnerability is well worth highlighting because it shows how extreme cloud identity risk can come from the plumbing—not just from a user clicking a bad link. The vulnerability was caused by older cloud components and unusual token flows that most admins never see day to day. Not only was it out of mind for most admins, but it was also out of sight, as the use of these tokens didn't generate logs in the victim tenant.  This would have made detection and response significantly harder if an attacker had found it first.

> The broader lesson here is straightforward and non-technical. Modern identity systems are built on trust boundaries, and legacy paths can quietly weaken those boundaries over time. It's also a reminder that the more identities with standing privilege you have, the greater the potential attack surface. CVE-2025-55241 reinforces the importance of treating identities and identity infrastructure as a critical attack surface, and the need to push for better visibility in the identity control plane so that unusual access quickly stands out.

# A Suite Deal for Attackers

Perhaps the most dramatic shift captured in this edition of the report is within productivity software. Love it or loathe it, the Microsoft Office suite is ubiquitous in the workplace, so any changes to its vulnerability landscape pose widespread implications.

Overall, Microsoft Office vulnerabilities surged by more than 3x year-over-year, rising from 47 to 157, with critical vulnerabilities jumping from 3 to 31—multiplying faster than an Excel spreadsheet—with a 10x increase from last year!

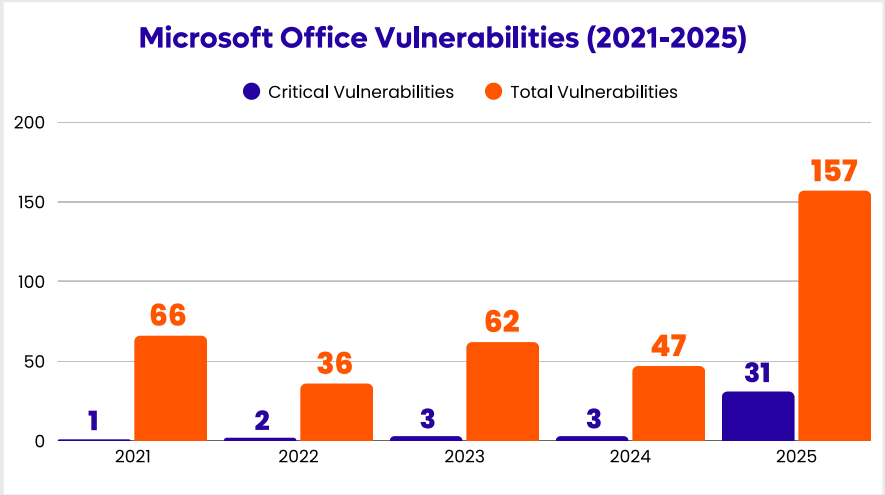


Figure 11: Microsoft Office total vulnerabilities increased by more than 3x year-over-year, with critical vulnerabilities growing by over 10x.

Microsoft Office is baked into many of our daily workflows and sits at the intersection of human behavior, daily operations, and business continuity. Macros, document sharing, preview panes, HTML rendering, and add-ins create a unique landscape for exploitation.

Not many other applications are required to do so many things on a daily basis for all types of roles and users. This fact also reminds us that when Office vulnerabilities spike, users remain the most reliable entry point as an attack vector via social engineering.

Microsoft has been on a long journey to harden Office. Today, many of the features that once allowed downloaded docs to freely execute macros and call OS functions have been heavily restricted by default. But Office's ubiquity still makes it a key focus for researchers and threat actors alike.

### The More-than-Remote Possibilities - CVE-2025-62557 & CVE-2025-62554

Looking at the examples of CVE-2025-62557 and CVE-2025-62554, both of these critical remote code execution vulnerabilities have a base CVSS score of 8.4. CVE-2025-62557 is a classic memory corruption vulnerability that impacted both desktop and mobile clients. CVE-2025-62554 is a "type confusion" where Office doesn't know how to handle an incompatible resource.

These critical vulnerabilities were combined to allow an attacker to craft malicious email attachments or links able to exploit the preview pane to execute code without requiring user interaction. Therefore, the blast radius of this attack vector is heavily dependent on the privileges of the exploited user, as the remote code execution occurs as the user. If the target has local administrator privileges, then so does the attacker, making it far easier to disable security controls and move laterally.

# AI, AI...oh!

There have been plenty of iconic wingmen in cinematic and literary history, but does Microsoft Copilot make the list? Well, it's a debatable question. Despite its boost to productivity, AI also brings new dimensions to vulnerabilities and the attack surface. Now three years after its launch, many organizations are continuing to uncover these Copilot risks retrospectively.

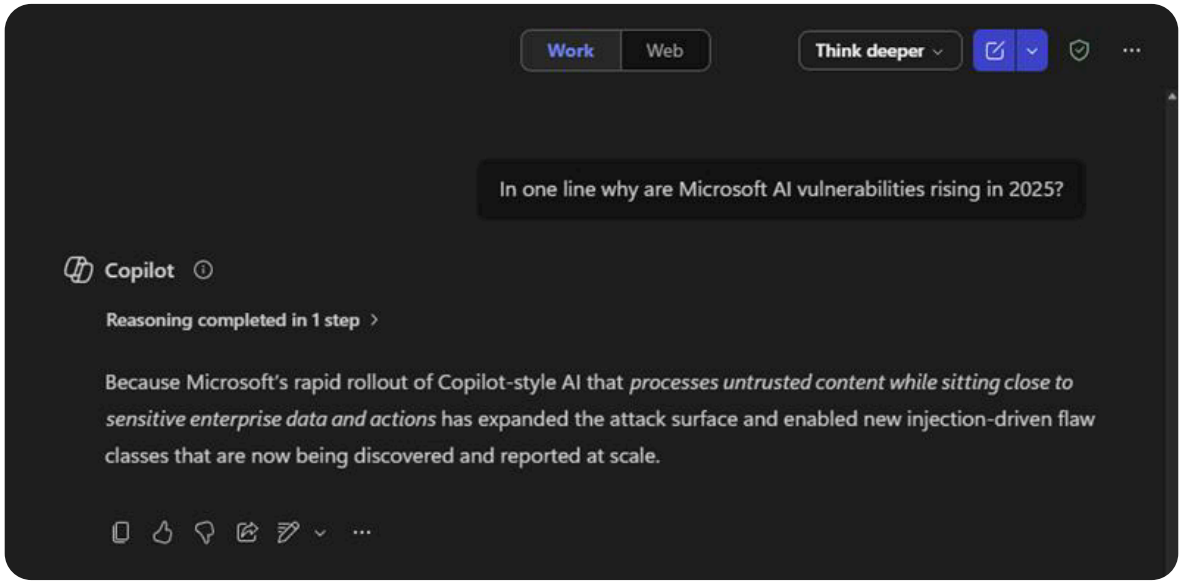


Figure 12: There are advantages to self-awareness

## Critical Microsoft Copilot CVEs of Note in 2025:

1. CVE-2025-32711 – EchoLeak: A critical vulnerability in Microsoft 365 Copilot caused by weak command input sanitization. It allows remote, zero-click exploitation to leak confidential data using AI model manipulation.

2. CVE-2025-59286: Allows unauthorized attackers to perform command injection and disclose sensitive information over a network command injection vulnerability in Copilot.

These vulnerabilities highlight how commonly deployed AI can provide an attacker with the ability to take a hijacked identity and exploit an AI agent and its privileged access using a malicious prompt. While you may no longer have a patch to deploy manually, you need to be aware of these vulnerabilities to inform your security and risk posture.

# We're Not Living on the Edge Anymore

We've come a long way from good ol' Internet Explorer, as this year, Microsoft Edge reached an all-time low of just 50 total vulnerabilities (0 of which were critical), representing an 83% decrease year-over-year. This is not accidental. Edge has benefited from rebasing to Chromium, which benefits from aggressive sandboxing, rapid update cycles, and strong isolation models. It proves that architectural and secure-by-design investments do deliver results. However, we should remember that browsers, even Chromium-based ones, are not inherently safe, and that reducing privilege and enforcing isolation pays long-term dividends.

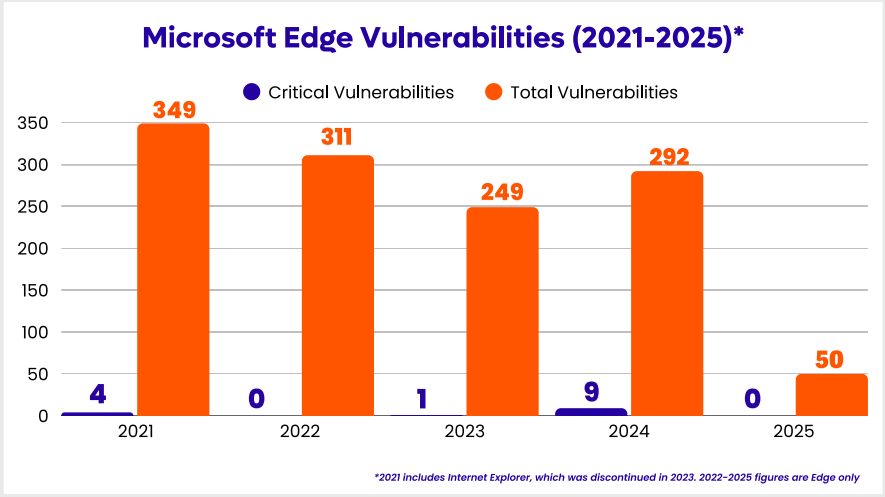


Figure 13: Microsoft Edge vulnerabilities reached an all-time low at 50 total vulnerabilities, representing an 83% decrease year-over-year—likely due to a number of secure-by-design investments.

In 2020, the year Microsoft rearchitected and rebased the code around Chromium, we saw 61 critical vulnerabilities in Edge. As the graph above shows, critical vulnerabilities then plummeted in ensuing years. While overall vulnerability numbers remained high, these were commonly low severity and would require chaining multiple separate exploits together to result in any real impact.

In 2025, we've seen overall vulnerability numbers drop dramatically for Edge. This is likely due to many factors, including some Chromium changes that focus on reducing user-abuse surfaces, restricting local network access, and improving memory safety. While there is no single dramatic change, the incremental strengthening of security, combined with massive market share, is resulting in significant real-world benefits.

# A Window into Overall Microsoft Risk

The Windows numbers continue to act like the canary in the Microsoft coal mine. Windows recorded 612 vulnerabilities in 2025, with 36 classified as critical, which is a modest rise from the prior year's 587 Windows vulnerabilities, with 33 critical in 2024. Taken together, this suggests the endpoint attack surface for Windows remains stubbornly broad. Windows is the place where "everything meets everything," spanning modern features and a long tail of backward-compatible components from 20+ years ago.

In practical terms, the data reinforces that Windows remains a reliable and high-value target for threat actors because it's where user activity, legacy services, and day-to-day operations converge into a large, diverse attack surface that is difficult to uniformly harden.

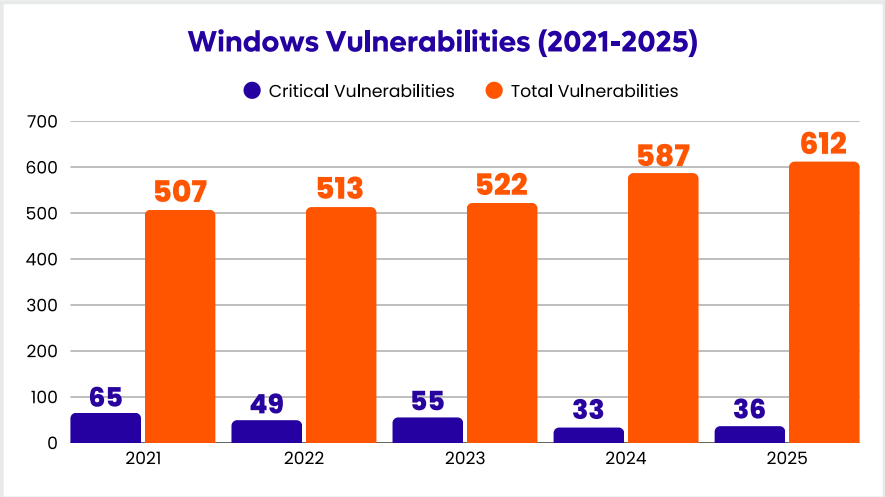


Figure 14: Windows vulnerabilities, both total and critical, increased slightly year-over-year.

### A Preview (Pane) of Things to Come

As the preview pane is built into Windows File Explorer and automatically loads content without the user having to directly open a file, it also serves as a tempting attack vector. This has clearly not gone unnoticed, as vulnerabilities that exploit the preview pane were a consistent theme throughout 2025.

### Critical vulnerabilities involving the Preview Pane:

CVE-2025-62554, CVE-2025-62557, CVE-2025-53731, CVE-2025-53740, CVE-2025-54910, CVE-2025-30377, CVE-2025-21298

The common mitigation advice across these CVEs is to disable the preview pane. As of October 2025 security updates, Microsoft now automatically disables the preview feature for files downloaded to mitigate a security vulnerability. With that said, this remains a trend worth keeping a close eye on. It's also worth repeating that implementing the principle of least privilege and application control on endpoints are key mitigations for these vulnerabilities, and the exploits for them.


Figure 15: The Preview Pane was a continuous source of critical vulnerabilities in 2025, allowing attackers to exploit file previews to gain code execution without the user opening the file.

# Serving Up Higher Stakes

On the server side, the story is more concerning from a strategic standpoint: Windows Server vulnerabilities climbed to 780 in 2025, with 50 critical, up from 684 vulnerabilities (43 critical) in 2024. This is a meaningful shift because servers are disproportionately "high-leverage" assets: they run shared services, form core infrastructure, and frequently operate with higher privilege and broader connectivity than most endpoints.

The data is effectively telling us that, while endpoint risk is persistent and wide, server-side risk is intensifying in a way that can amplify impact. When server vulnerabilities turn critical, the potential for rapid disruption, lateral movement, and an infrastructure-wide blast radius becomes the defining problem to manage.

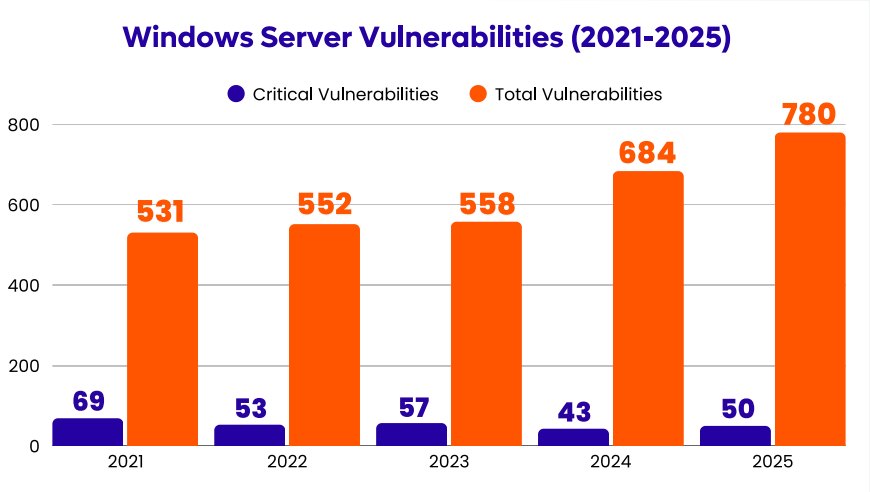


Figure 16: Windows Server vulnerabilities showed a rising concentration of high-impact issues, with total vulnerabilities increasing from 684 to 780, and critical vulnerabilities also increasing.

The shared takeaway from Windows and Windows Server this year is that “volume” alone is less informative than where impact concentrates and how privilege magnifies it. Windows Desktop shows persistent exposure at scale, while Windows Server shows a rising concentration of high-impact issues in the systems most likely to underpin business-critical operations.

In report terms, this supports a prioritization narrative: treat critical and high-severity Windows Server fixes as time-sensitive because they can quickly collapse trust boundaries. Also pair patching with compensating controls that reduce blast radius, especially tighter privilege controls and segmentation around server roles. The potential for compromise is always most dangerous where privilege and shared infrastructure intersect.

# Vulnerabilities by Category

## CTRL+C, CTRL+V

In 2025, the overall category distribution of vulnerabilities essentially remained copy-paste from previous years. Elevation of Privilege (EoP) and Remote Code Execution (RCE) continue to dominate the landscape.

It's a pattern we see year after year in this report, and it's clear as to why: EoP and RCE represent the two primary goals of any threat actor looking to exploit a system. In short, attackers aren't just seeking access; they're seeking the power to execute and move within a system. These two attack paths offer them the level of privilege they need to achieve their objectives.

In terms of historic trends, RCE has fallen back in line with previous years after a spike in 2024, and EoP is slightly under the five-year trend average, but not quite down to the level of 490 we saw in 2023.

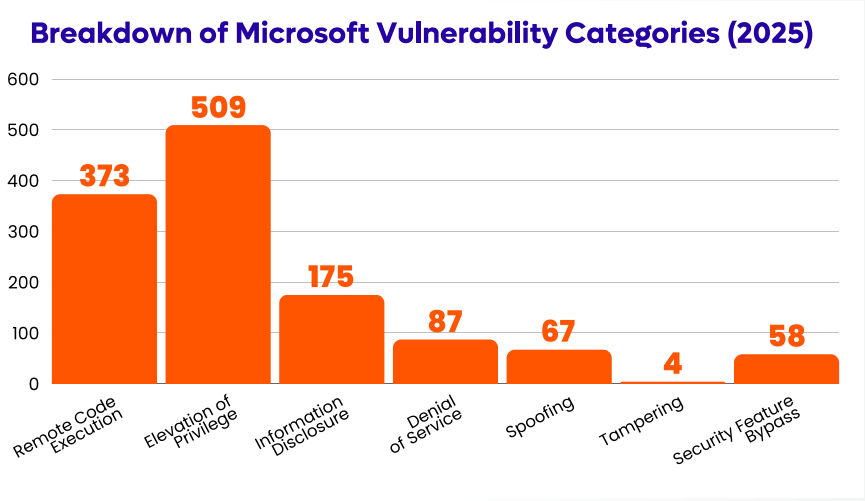


Figure 17: As seen in previous years, RCE and EoP remained top vulnerability categories. Additionally, Information Disclosure showed a significant jump from 101 to 175.

## Leaks can Become Floods

Information Disclosure vulnerabilities leapt from 101 in 2024 to 175 in 2025. These are generally classified as less severe than other vulnerability categories as they are "just some information leakage". However, if you unintentionally leak any type of sensitive information or technical details, it can help an attacker understand and map your environment.

As the Microsoft estate gets broader and more interconnected, there are more places from which information can leak. The more that leaks out, the greater the risk of that data becoming a pathway into your environment.

This trend can indicate a greater interest in reconnaissance as attackers seek to better understand the broader Microsoft estate and evolve their techniques.

# Guardrails or On-Ramps?

Historically, we've observed attackers use security 'guardrails'—mostly outdated controls—as on-ramps into Microsoft ecosystems. Yet, we see a trend reversal for Security Feature Bypass vulnerabilities. The number of vulnerabilities in this category tripled from 30 in 2020 to 90 in 2024, but in 2025, it dropped 36%, to 58.

The previous spike occurred because attackers targeted historic security controls, like the Mark of the Web: a tag added to files to indicate they originated from the internet rather than the local network, signaling they should not be trusted. These long-lived security controls were designed for a different era and made for easy targets. But after a period of playing whack-a-mole, Microsoft has hardened these security features and introduced newer controls that are more resistant to exploitation.

However, we should still expect to see targeting of these security guardrails as alternatives to classic bug exploitations. There is a lot of creative effort invested in sleuthing ways to work around key security features, so it will be interesting to see if the downward trend persists.

# What does 2026 hold?

AI revolution, robot uprising, or more of the same?

This year has already started with a bang. The first Patch Tuesday of 2026 landed with security updates for 114 flaws, including three zero-day vulnerabilities, with one reported as actively exploited. That is a strong reminder that we start the year on the back foot if we treat patching as a once-a-month routine instead of a continuous discipline.

But the bigger story for 2026 isn't just how quickly we patch—it's what happens when a vulnerability gives an attacker a foothold and they immediately pivot into exploiting identities and privilege. Securing identity is as pivotal now as timely patching, especially because vulnerable applications are increasingly operated by privileged non-human identities, including AI agents that can spin up and down quickly in cloud environments.

So, where does the "ghost in the machine" come in? Vulnerabilities don't exist in a vacuum—they're exploited through identities, and the identity landscape is undergoing the largest expansion in history as machine and agentic AI identity growth outpaces humans by orders of magnitude.

This expanded identity attack surface is what BeyondTrust Phantom Labs™ calls the ghost in the machine: non-human identities (NHIs) such as service accounts, APIs, containers, bots, and agents. Designed to authenticate and operate without human involvement, these identities frequently lack defined ownership and oversight. They power automation and scale across cloud, DevOps, SaaS, and AI, making them indispensable to modern systems.

The same properties that make non-human identities useful—always-on execution, broad connectivity, and privileged access—also make them attractive targets. NHIs often lack MFA equivalents, rely on long-lived credentials and secrets, are rarely monitored, and commonly suffer from excessive privilege and weak lifecycle discipline.

So alongside software vulnerabilities, we have a parallel risk of identity misconfigurations, compromise and abuse. These can create new kinds of vulnerabilities: ones that don't get assigned CVEs, but can have the same critical consequences.

When identity plus privilege becomes capability, reducing standing privilege is one of the most reliable ways to shrink the blast radius when zero days come knocking.

# Expert Commentaries

## Sami Laiho

**Senior Technical Fellow at Adminize, & Microsoft MVP**

The findings in the 2026 Microsoft Vulnerabilities Report reinforce a reality that many security practitioners have understood for years, but organizations are still struggling to operationalize: vulnerabilities do not create impact—privilege does.

While the slight reduction in total vulnerabilities may appear encouraging at first glance, the doubling of critical vulnerabilities tells a far more important story. The data shows a clear concentration of risk in vulnerabilities that enable Elevation of Privilege (EoP), which continues to dominate the landscape. This is not coincidental. It reflects the fundamental attacker playbook: initial access is rarely the end goal—privilege escalation is.

From a Principle of Least Privilege (PoLP) standpoint, this trend should not be viewed purely as a software security problem, but as a systemic design issue. The continued success of EoP exploits indicates that environments still rely heavily on excessive standing privileges, broad administrative roles, and weak separation of duties. In such environments, even low-complexity vulnerabilities can quickly translate into full domain compromise or control over cloud tenants.

The report also highlights a critical evolution in the attack surface: identity has become the primary control plane. Vulnerabilities in Azure, Entra ID, and cloud-integrated services demonstrate that modern attacks are no longer confined to endpoints or servers—they target the systems that define trust itself. When identity systems are compromised, attackers do not need to "break in" further; they can simply operate as trusted entities. This is particularly concerning given the rise of non-human identities, service principals, and AI-driven agents, which often operate with persistent and excessive privileges that are rarely scrutinized with the same rigor as human accounts.

Another key takeaway is the role of privilege in determining the real-world impact of vulnerabilities. Many of the critical vulnerabilities described—particularly in Microsoft Office and preview-based attack vectors—execute in the context of the user. The difference between a contained incident and a widespread breach is therefore not the vulnerability itself, but whether the compromised identity has administrative rights. This is where PoLP delivers its most tangible value: it directly limits the blast radius of successful exploitation.

The report correctly emphasizes that patching remains essential, but it is no longer sufficient. In a landscape where zero days, identity abuse, and AI-driven attack paths are increasingly common, organizations must assume that compromise will occur. The question is no longer "can we prevent exploitation?" but rather, "what happens when it does?" Least privilege is one of the few controls that consistently reduces impact across all attack vectors—known and unknown.

From a strategic perspective, this means shifting focus from reactive vulnerability management to proactive privilege reduction. This shift includes eliminating standing administrative rights, enforcing just-in-time access models, tightly controlling Tier 0 assets, and extending least privilege principles to service accounts, APIs, and AI agents. In cloud environments where privilege can scale instantly and globally, this becomes even more critical.

Ultimately, the report validates a core security principle: the true risk in modern environments is not the presence of vulnerabilities, but the presence of unnecessary privilege. Organizations that continue to prioritize patching without addressing privilege will find themselves repeatedly exposed to the same attack patterns. Those that embrace least privilege as a foundational design principle will not eliminate vulnerabilities, but will dramatically reduce their ability to cause harm.

# Paula Januszkiewicz

**CEO of CQURE Inc. and CQURE Academy, Security Expert, Penetration Tester and Trainer, Microsoft MVP on Security and Microsoft Regional Director**

In 2026, attacks are not just incidents. They are consequences of architectural weaknesses.

## The Illusion of Sophisticated Attackers

There is a persistent belief in cybersecurity that organizations lose because attackers are becoming more advanced, more creative, or more sophisticated. In reality, the opposite is often true. Most organizations do not fail because of groundbreaking techniques, but because their environments allow attacks to succeed. This is a pattern consistently observed in real-world incidents and reinforced by the latest Microsoft Vulnerabilities Report, where the total number of vulnerabilities appears relatively stable, yet the number of critical vulnerabilities has doubled. This shift is not just statistical. It reflects a deeper problem in how environments are designed, managed, and understood.

## The Real Beginning of an Attack

The key issue is not where the attack starts, but what happens next. Organizations often treat security as a matter of preventing entry, focusing heavily on perimeter defenses, patching, or individual vulnerabilities. However, the initial access point is rarely the decisive factor. It is merely the starting condition. What truly defines the success of an attack is how the environment responds to that initial foothold. From an attacker's perspective, the process is rarely chaotic. It follows the logic of the infrastructure itself. Relationships between systems, identities, and privileges create predictable paths. These paths are not invented by attackers. They already exist, embedded in the architecture, waiting to be used.

## Misconfiguration: The Hidden Root Cause

This is why misconfiguration remains one of the most critical and underestimated risks in modern cybersecurity. Vulnerabilities such as remote code execution or privilege escalation may trigger an incident, but they do not determine its outcome. The outcome is shaped by excessive privileges, implicit trust relationships, lack of segmentation, and poor visibility. In many environments, permissions are granted incrementally over time, without full awareness of their cumulative effect. Identities become overprivileged, systems become interconnected in unintended ways, and boundaries blur. In such conditions, even simple attack techniques become highly effective. Credential theft, Pass-the-Hash, and Kerberoasting do not rely on innovation. They rely on consistency and opportunity.

This is where the need for upskilling becomes critical. Organizations must move beyond tool-centric thinking and develop a deeper understanding of how attacks actually unfold. Technology alone cannot compensate for a lack of awareness. Security teams need to understand not only how systems are built, but how they behave under stress, how they fail, and how they can be abused. Without this perspective, even well-equipped organizations remain reactive, addressing symptoms rather than causes.

### Beyond a Single Vendor Perspective

It is also important to recognize that risk does not reside within a single vendor ecosystem. While Microsoft technologies play a central role in many infrastructures, they are only part of a broader landscape. Recent vulnerabilities affecting edge devices clearly demonstrate this. Well-known vulnerabilities from 2024 are still being actively exploited today, allowing attackers to bypass perimeter defenses and gain direct access to internal environments. What makes such vulnerabilities particularly dangerous is their position within the architecture. Edge systems define trust boundaries. When they are compromised, the distinction between internal and external collapses. This highlights a fundamental principle: cybersecurity is not about securing individual technologies, but about securing the relationships between them.

### The Limits of CVE Thinking

Another limitation in how organizations approach risk is the reliance on CVEs as the primary measure of exposure. While CVEs provide valuable insight into known software flaws, they do not capture the full reality of modern attack surfaces. Many critical attack paths do not originate from software vulnerabilities at all. Instead, they emerge from identity abuse, token misuse, excessive permissions, or poorly governed automation. These risks are often invisible in traditional vulnerability management processes, yet they play a decisive role in incidents.

### AI Agents and the Expanding Identity Surface

These challenges are becoming even more pronounced with the rise of AI and automation. Modern environments are no longer composed solely of human users. They include service accounts, APIs, automation workflows, and increasingly, AI agents. These entities authenticate, access resources, and execute actions, often with significant privileges. However, they are rarely managed with the same level of control as human identities. They lack ownership, visibility, and lifecycle governance. As a result, they introduce a new category of risk, one that does not fit neatly into existing security models. An attacker who compromises such an identity does not need to escalate privileges, because the privileges are already there.

### Transparency as a Control Layer

In this context, transparency becomes one of the most critical elements of effective cybersecurity. During incidents, organizations frequently struggle to answer three fundamental questions: what happened, where it spread, and what was accessed. When these questions

cannot be answered, response efforts become fragmented and ineffective. The problem is not necessarily the absence of tools, but the absence of visibility. Telemetry is often dispersed across systems, logging is inconsistent, and critical context is missing or delayed. In such conditions, organizations are forced to reconstruct events instead of managing them in real time. As emphasized in practice, transparency is not about reporting. It is about control. It transforms uncertainty into evidence, and evidence into action.

### Least Privilege as a Strategic Defense

Despite the complexity of modern environments, some principles remain consistently effective. Least privilege is one of them. It is often discussed, but rarely implemented comprehensively. When properly applied, it significantly reduces the potential impact of an attack. If an attacker gains access to a low-privileged identity, their ability to move within the environment is limited. If the compromised identity holds excessive privileges, the attack can escalate immediately. Least privilege does not eliminate risk, but it constrains it. It reduces the blast radius and creates opportunities for detection and response. Importantly, this principle must extend beyond human users to include services, applications, and AI-driven processes.

### Cybersecurity as a Layered System

Ultimately, cybersecurity cannot rely on a single control or approach. It must be understood as a layered system. Patching addresses known vulnerabilities, but does not prevent misuse of valid access. Detection enables response, but often after the attack has already progressed. Identity security limits privilege, but must be supported by segmentation and monitoring. Each layer compensates for the limitations of the others. This layered approach reflects a fundamental reality: vulnerabilities will continue to exist, exploits will continue to emerge, and some level of compromise is inevitable.

### Conclusion: Control Over Perfection

The objective, therefore, is not to eliminate all risk, but to maintain control over it. Organizations that succeed are not those that avoid incidents entirely, but those that can effectively detect, understand, and contain them. This requires a shift in mindset, from focusing on individual threats to understanding the broader system in which those threats operate.

In the end, cybersecurity is not defined by tools or technologies, but by the quality of decisions made about architecture, identity, and visibility. The attack itself is rarely the core problem. It is the natural consequence of an environment that was not designed to withstand it.

# Jane Frankland, MBE

**Founder of the IN Security Movement,**
**CEO of KnewStart & Best-Selling Author**

530 vulnerabilities in 2016. 1,273 in 2025. That near-tripling growth over a decade isn't a trend line; it's a structural shift in the attack surface. While 2025's headline figure represents a slight dip from 2024's peak, the detail underneath tells a more urgent story. Critical vulnerabilities didn't dip. They doubled. In cyber, volume is noise, severity is the signal, and that signal is loud.

What the data is telling us is that Elevation of Privilege (EoP) dominated, followed by Remote Code Execution. Together they represent the most dangerous combination in an attacker's toolkit—get in, then escalate. EoP figures are particularly telling as they're not just technical weaknesses; they're the consequence of organisations that haven't implemented least privilege properly.

The one vulnerability category that didn't fall deserves its own moment. Information Disclosure was the only category to move in the wrong direction and it's often the quiet precursor to everything else. What an attacker learns about your environment, your architecture, and your users is what makes the next stage possible. In a landscape where AI is being used to personalise attacks and competitive espionage is increasingly digital, data that leaks silently arms your adversary.

The product picture offers good and bad news. Microsoft Edge is a rare good news story. Total vulnerabilities fell dramatically, with critical vulnerabilities dropping to zero. We see that focused investment can shift the trajectory. The question is where else that discipline gets applied.

Unfortunately, both total and critical vulnerability counts in Windows Server continued to rise. Microsoft Office more than tripled in total vulnerabilities, with critical vulnerabilities jumping tenfold. Office is where your people live—email, documents, collaboration. So when the human layer and the technical layer converge, risk compounds.

The figure that stopped me is Azure and Dynamics 365. Total vulnerabilities held steady, but critical vulnerabilities increased nearly tenfold in a single year. A critical vulnerability in Azure isn't contained to Azure—it's a threat to identity, data, connectivity, and everything built on top of it.

One more note on measurement. Microsoft's own severity rating recorded 157 critical vulnerabilities in 2025. The National Vulnerability Database (NVD) scored only 42. If your patching decisions are driven exclusively by NVD scores, you may be significantly underestimating your exposure. Use both. Understand the gap.

What to do about it? Patch quickly for what's exploited in the wild. Patch consistently for everything else. Never let "we'll get to it" become a six-month gap, and don't assume the cloud patches itself!

Apply least privilege rigorously. Every user, every service account, every application should operate with only the access it genuinely needs. This limits the blast radius when something is exploited.

Monitor inside the environment, not just at the perimeter. The information disclosure trend means data is leaving through vulnerabilities that often trigger no alerts.

Finally, a word on AI agents.

The data doesn't yet capture this, but it will. AI agents inherit identity, access, and privilege. Most are being deployed without the governance rigour we apply to human accounts. The Azure critical vulnerability spike matters here as this is the infrastructure layer where AI services live, authenticate, and interact with your data. A near-tenfold increase in critical vulnerabilities in that environment, combined with ungoverned machine identities operating autonomously within it, is a converging risk, not a theoretical one.

Apply the same least privilege principles to your AI agents that you apply to your people. Govern their identity, monitor their behaviour. Understand what they can access and what they can expose too.

The vulnerabilities are documented. The patches exist. The principles are well understood. The question, as always, is whether we act before or after the incident.

# Katie Moussouris

**Founder and CEO**
**Luta Security**

### Pay No Attention to the Numbers in Front of the Curtain

The 6% drop in total Microsoft vulnerabilities is the wrong number to watch. Critical vulnerabilities doubled, from 78 to 157. Azure / Dynamics 365 saw a 9x increase in critical vulnerabilities, and Microsoft Office saw a 10x increase.

### CVEs Don't Cover the Whole Picture

Many of the most consequential cloud and AI vulnerabilities never receive CVEs. CVE-2025-55241 let an attacker forge tokens accepted across any Entra ID tenant, left no logs, and came from legacy infrastructure most admins never think about. CVSS scores also only reflect what is known at patch time. They do not replace knowing your own environment.

### Least Privilege Remains the Most Reliable Control

Removing local admin rights has historically mitigated around 75% of Microsoft's critical vulnerabilities. The seven preview pane CVEs from 2025 make this point: zero-click exploitation against a standard user is a nuisance; but against a local admin, it is a full compromise. Non-human and cloud identities rarely receive equivalent scrutiny, but should.

### The AI Agent Surface Is Already Here

EchoLeak (CVE-2025-32711) enabled zero-click remote exploitation through Copilot model manipulation with no user interaction required. AI agents hold persistent privileged access, run without human oversight, and can be manipulated in ways traditional controls miss. Most organizations would not give a new employee standing admin access with no monitoring and no lifecycle governance. That is exactly what they are doing with agentic AI today.

# David (DJ) Morimanno

**Field CTO**
**Xalient**

What stands out to me in this report is not that Microsoft's total vulnerability count dipped slightly in 2025. It is that the risk became more dangerous, more concentrated, and more aligned to the systems that matter most. Total vulnerabilities fell, but critical vulnerabilities doubled. That is the real signal. This is not a story about volume. It is a story about severity, exposure, and control.

This brings us back to a point security leaders still need to internalize: identity and privilege remain the real battleground. Attackers do not win simply by landing on a box. They win when they can impersonate, elevate, move laterally, and gain trusted access to the next system. That is why Elevation of Privilege continues to dominate this landscape. It is also why Microsoft's own 2025 Digital Defense Report is so telling. Microsoft reports that the vast majority of identity attacks were password spray attacks, and it specifically warns that adversaries are increasingly targeting workloads such as apps, services, and scripts because these non-user identities often carry high privilege with weak controls.

This is also where the ghost in the machine becomes real. Non-human identities are no longer a side issue. Service accounts, automation, API-driven workloads, service principals, and now AI agents are becoming active participants in business operations. The problem is that many organizations still govern them like background plumbing instead of treating them as identities with access, authority, and risk. Microsoft is explicit that AI introduces new attack surfaces, including prompt injection, malicious tool invocation, and poisoned data. That means the enterprise is not just managing users anymore. It is managing human identities, machine identities, and, increasingly, autonomous actors making decisions inside trusted systems.

That is why patching alone is not enough. Patching is essential, but it is not a security strategy by itself. If privilege is excessive, trust is assumed, and access paths remain overexposed, then a patched environment can still be an insecure environment. We saw this clearly in Microsoft's response to the active exploitation of on-premises SharePoint vulnerabilities in July 2025. The issue was not just the flaw. It was the combination of exposed infrastructure, trusted application context, and the attacker's ability to execute, persist, and steal keys.

The ultimate factor here is trust. Trust in the identity. Trust in the device. Trust in the workload. Trust in the agent. Zero Trust matters because modern defense is no longer about assuming trust and then reacting. It is about continuously validating trust, constraining privilege, and governing every identity (human and non-human) as if it can become the next attack path. That is the lesson this report should leave with every security leader.

# BeyondTrust Phantom Labs™

## Research Team

Phantom Labs™, BeyondTrust's dedicated research team, applies a "think like attackers" perspective to accelerate identity security innovation and deliver actionable insights for defenders.

From our viewpoint, vulnerabilities don't exist in a vacuum, or in a disclosure framework. Vulnerabilities are a means to an end to exploit privileges. And thanks to legacy cloud infrastructure bearing the weight of rushed agentic AI solutions, the identity threat landscape is undergoing the largest expansion in history, further multiplying these types of identity risks.

This explosion of non-human identities is what we frame as a "ghost in the machine": NHIs in the form of service accounts, APIs, containers, bots, certificates, and agents. These NHIs authenticate and act without human interaction, often lacking clear ownership or governance. They exist for automation, integration, and scalability, and modern cloud, DevOps, SaaS, and AI systems cannot function without them.

The AI security market is racing toward runtime defenses, red-team automation, and policy toolkits. Yet, none of them solve the root problem: autonomous actors operating on identity systems never designed for them.

The Phantom Labs team's hot take? Microsoft's CVSS and vulnerability disclosures can't fix the agentic AI house of cards. Attackers are exploiting vulnerabilities in agentic AI systems at an accelerated rate, and customers are vulnerable.

Many organizations are unaware of where these agents exist, let alone what level of access they have and which actively-exploited vulnerabilities they are exposed to. It's not enough to monitor and respond to CVSS disclosures; they don't disclose every risk and can be months delayed in resolving the active exploitation.

Over the past year, **Phantom Labs™** has observed a 466.7% increase in AI agents operating within enterprise environments.

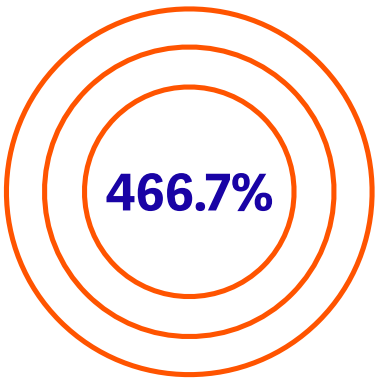


This is why the Phantom Labs research team is dedicated to agentic AI security research, leading the charge on how the industry can be successful in safely adopting the full potential of AI.

Learn more about our most recent research on agentic AI and beyond:

Pwning AI Code Interpreters in AWS Bedrock AgentCore

How Command Injection Vulnerability in OpenAI Codex Leads to GitHub Token Compromise

Entra ID App Escalations: Attacks & Defenses

Learn more about Phantom Labs™ and how we "think like attackers" to expose new privilege escalation paths and identity attack vectors.

# Bradley Smith

**SVP and Deputy CISO**
**BeyondTrust**

The data in this report should be uncomfortable—not because the numbers are surprising, but because the pattern they reveal is one the industry has been choosing to ignore.

Critical vulnerabilities doubled year over year. Elevation of Privilege continues to dominate at 40% of all vulnerabilities. Cloud identity flaws like CVE-2025-55241 demonstrate that a single compromised identity can collapse trust boundaries across an entire tenant. And yet, the dominant response from most organizations remains the same: patch when we can, budget when we must, act when something breaks. We have to stop betting that the breach will happen to someone else first.

My primary role is leading the defense of BeyondTrust and its customers, many of whom operate across the mid-market and critical infrastructure supply chains. Every day, I see an exploitation timeline that has fundamentally changed.

At the same time, nation-state actors aligned with adversaries like Iran are actively targeting identity infrastructure, exploiting the same identity attack vectors this report highlights as the dominant vulnerability category. The exploitation window is compressing from both directions: AI accelerates weaponization while geopolitical escalation expands the adversary set. This data is not academic. It describes the attack surface that real threat actors are operating against right now.

The gap has never been knowledge. It has always been the willingness to act on what we already know, before the incident forces the decision.

That requires something this industry talks about far less than it should: courageous leadership. Not recklessness, not blank checks, but the deliberate decision to invest in security architecture, least privilege enforcement, and identity governance before the breach justifies the budget. It means boards that treat cybersecurity as a strategic function rather than a compliance checkbox, and executives who fund security programs based on the threat landscape the organization actually faces.

The ghost in the machine that this report describes is real. But the more dangerous ghost is the one in the boardroom: the assumption that inaction is the safe choice. It never was, and the data in this report makes that harder to ignore than ever.

The question is not whether your organization will face a critical vulnerability. It is whether your leadership will have made the investments that determine how that story ends.

# Kevin E. Greene

**Public Sector Chief Cybersecurity Technologist**
**BeyondTrust**

Almost 75% of CVEs in CISA KEV are privilege-dependent, meaning these CVEs require existing privileges to exploit, or they grant and accelerate privileged access as an outcome. Once exploited, these CVEs will produce meaningful progression and impact. Threat actors can only weaponize CVEs that are compounded in nature, meaning when Remote Code Execution fires, it will need to inherit some form of privilege on the target to be successful. The truth of the matter is that a CVE unable to exploit existing privileges or grant and accelerate privileged access is operationally ineffective, even if it has a CVSS score of 10.

This is why a Privilege Disruption architecture is needed to put weaponized CVEs into a chokehold, preventing and denying threat actors from converting their initial access into meaningful control through a privilege plane.

This is increasingly important for the ghost in the machine highlighted in this report. Non-human identities and AI agents must be monitored, tracked, and governed with least privilege in mind. They become attack vectors through silent identity abuse, which gives these CVEs the privilege planes to operate and be successful. Understanding the privilege scope and reach these identities hold is essential, which is why a "Leave No Privilege Behind" mindset is needed to zero in on where privilege debt accumulates across the identity estate.

This report is a reminder that all software has vulnerabilities and weaknesses—which is why removing EoP and RCE vulnerability classes through secure-by-design activities, and disrupting them when they appear, is essential to shrinking the window of exposure for weaponized CVEs. Patching has always been an Achilles' heel; the subtle but important dominance of RCE and EoP in this report signals that successful exploitation of these CVEs will result in progression and impact. I believe these CVEs matter the most because in cases where they are compounded, as seen with CVE-2025-32711, CVE-2025-21298 and CVE-2025-65037, they enable and accelerate privilege escalation and abuse, converting initial access into mission capability in a single exploit chain. Treating CVE triage through a Privilege Disruption framework makes sense given that threat actors think in capability gaps and only need one path to privilege.

In the end, we must operationalize Leave No Privilege Behind and Privilege Disruption to put these CVEs in a chokehold before they can wreak havoc on our critical infrastructure systems, mission capabilities, and business functions. This does not absolve organizations from patching—it gives them the breathing room to better manage risk and the operational agility to be responsive to evolving threats. More importantly, it imposes cost, uncertainty, and risk on adversary operations, forcing threat actors to burn through their CVE arsenal against an architecture that denies them the privilege control they need to succeed.

# Marc Maiffret

**Chief Technology Officer**
**BeyondTrust**

I've been finding and disclosing Microsoft vulnerabilities since the late 1990s. Back then, the attack chain was straightforward: get a foothold, exploit a buffer overflow, get SYSTEM. Privilege escalation was almost automatic. The hard part was the initial foothold.

Twenty-five years later, the hard part has shifted. But privilege escalation? It never left. Elevation of Privilege accounts for 40% of all Microsoft vulnerabilities. Place that alongside identity sprawl and over-privileged accounts, and you can see exactly why threat actors are so focused on exploiting identities. The doubling of critical vulnerabilities year-over-year, from 78 to 157, reinforces the point. This isn't a blip; it's a signal that we shouldn't get complacent as the risks are very real.

Nowhere is this clearer than in the cloud. CVE-2025-55241, a flaw in Entra ID token handling, would have allowed any attacker to impersonate any identity, including Global Administrators, across any tenant, with no user interaction and no logs in the victim tenant. This is exactly the kind of vulnerability that redefines what "critical" means.

When I started in this industry, that level of access required being on the network. Today, one compromised identity or one cloud vulnerability can collapse trust boundaries at machine speed.

This is why shifting our perspective matters. CVE counts have always been an incomplete picture. Identity misconfigurations, over-privileged machine accounts, AI agents with unconstrained access: these don't get CVEs, but they have the same critical consequences. Instead, we need to connect the data to how attacks actually happen. “Thinking like an attacker” is the lens security leaders need to be operating through (and it’s one that our Phantom Labs team is continually focusing on).

Attackers aren't trying to overpower your security stack. They're looking for an over-privileged service account, the unpatched critical vulnerability, and the AI agent with unconstrained access. One good position and the submission is inevitable.

That is why the “ghost in the machine” reference in the report touches on something very real. Non-human identities are expanding faster than governance can track, carrying high privilege without scrutiny, MFA, or behavioral baselines. These are just the positions an attacker can leverage. If I were building an attack chain today, I wouldn't start with a kernel exploit. I'd start with an over-privileged service account in a cloud environment with no monitoring.

The ghost in the machine is real. The only question is whether you see it before the attacker does.

# Best Practices for Mitigating Microsoft Ecosystem Risks

We've witnessed a dynamic evolution in the Microsoft landscape over the 13 years of this report. The ecosystem has expanded massively, and while some challenges such as security bypass attacks and Edge vulnerabilities have been reduced, others, such as Copilot risks and Office vulnerabilities, have been introduced or expanded. But all the same, the future remains unknown, despite our most educated guesses.

With that said, the following security practices are a very good bet to improve security and cyber resilience success for your Microsoft environment and other ecosystems.

## 1. Tailor Vulnerability Management to Your Environment

A one-size-fits-all approach to patching and vulnerability management doesn't work. While patching sooner rather than later can help prevent a vulnerability from metastasizing into a major security incident, you must also consider the patch's impact on the environment, and the potential operational risks it may introduce.

Without the context of your own organization's threat models, you can't fully understand how best to prioritize patches and / or use other security hardening controls and mitigations. What is considered a critical patch for one organization may not be true for another. Leverage your business context to create the vulnerability management strategy that best addresses your risks, but also leverage tools that can help you properly gauge business impact.

With that said, ensuring your operating system and third-party software are up-to-date and avoiding the use of end-of-life software in your environment remain essential for good cybersecurity hygiene.

## 2. Implement Least Privilege and Zero Trust Controls Across the Stack

As we continue to see the emergence of new technologies, attack techniques, and risks, a multilayered least privilege approach stands as a powerful defense for the Microsoft ecosystem and beyond.

Least privilege is a proactive approach, offering a layered defense for the places where patching might fail or might not be available in time. It also puts an adaptive strategy in place to safeguard the explosion of NHIs, including agentic AI, that are increasingly pervasive in enterprise environments.

The more cohesive your least privilege approach across the network, identities, accounts, endpoints, applications, sessions, clouds, on-premises environments, and so on, the fewer the gaps that give threats a foothold or allow them to execute lateral movement and escalate.

As this report has highlighted in past years (when more specific privilege-level data was available from Microsoft), removing local admin rights alone and controlling execution has historically mitigated 75% of Microsoft's critical vulnerabilities.

A strong, privilege-centric identity security strategy is essential for discovering, controlling, and managing directly assigned entitlements and privileges, as well as potential escalation paths. By proactively hardening all types of identities—human, machine, and agentic AI—this least-privilege approach effectively reduces the threat surface and limits the blast radius of attacks, even when patches are not yet applied.

## 3. Secure Remote Access Pathways

Traditional remote access technologies, such as Microsoft's Remote Desktop Protocol (RDP) and VPNs, continue to provide common entry points for attacks. These technologies are often stretched beyond their proper use cases, resulting in security exposures and breaches.

To ensure access pathways are hardened and secure, adhere to these best practices:

- Ensure RDP is not exposed to the internet.
- Never allow VPNs and BYOD to mix.
- Enforce strong authentication and session monitoring to detect misuse.
- Replace VPNs or augment them with zero trust security controls for vendor access and privileged access use cases.
- Enforce least privilege and just-in-time access.

## 4. Implement Identity Threat Detection and Response (ITDR)

Organizations need to make proactive security posture changes and corrections to prevent threats from gaining a foothold, and they also need to orchestrate a rapid response to prevent or mitigate any damage from emerging attacks.

Identity Threat Detection and Response (ITDR) is a multi-discipline approach that aims to integrate capabilities for holistic identity security visibility, threat detection, investigation, and response. This begins with complete visibility and observability into the True Privilege™ of identities and the identity security posture across all enterprise identity stores (Active

Directory, Entra ID, Okta, Ping, and more). From there, organizations can gain understanding of the escalation pathways that can operate through those identities, and which steps are needed to improve posture or break an attack.

## 5. Prepare for Tomorrow

Understanding Microsoft vulnerability trends goes a long way toward making more informed decisions, keeping your organization secure, and managing emerging threats. The past few years have ushered in some security improvements, but have also laid bare for us just how dizzying the speed of change can be.

We're no longer just dealing with endpoints, but a vast hybrid environment of on-prem, cloud, SaaS, and OT, alongside an explosion of automated processes, driven increasingly by machines and AI agents. As such, we're likely to see many impactful shifts in threats, and at a faster pace.

As an industry, we're getting better at patching in a timely manner, but the vulnerabilities and patches are coming at us faster and in higher volumes than ever before. What's more, armed with AI-powered tools, attackers can find and exploit vulnerabilities faster than ever.

But as the saying goes, the best offense is a good defense. By reducing software and identity vulnerability exposures, you reduce the risk an attacker can move laterally and execute privilege, let alone gain a foothold in the first place.

# Privilege-Centric Identity Security from BeyondTrust

BeyondTrust enables organizations to control privileges and, as a result, reduce risks throughout their environments. We provide a privilege-centric identity security solution set designed to move customers from simply managing access, to controlling privilege everywhere—IT, OT, humans, machines, AI agents, etc. Our Pathfinder Platform integrates our entire solution set, and unifies visibility, intelligence, and protection within a single console.

Today, our multicategory identity security leadership spans Privileged Access Management (PAM), Identity Threat Detection and Response (ITDR), Secrets Management, and Cloud Infrastructure Entitlement Management (CIEM). Pathfinder unites all these disciplines, enabling customers to better secure every identity, everywhere.

**Customers rely on BeyondTrust to:**

- Gain cross-domain visibility and understanding of their entire identity security posture, including the True Privilege™ of every identity—human, machine, and agentic AI.
- Visualize entitlements and Paths to Privilege™, including those that other solutions miss.
- Implement a true least privilege model that removes admin rights and standing access, consistent with zero trust principles.
- Secure remote access pathways and infrastructure by ensuring all access—whether by human, machine, employee, or third-party / vendor—is granularly controlled and audited.
- Prevent account hijacking and privilege escalation by securely managing all human and NHI privileged credentials, DevOps secrets, SSH keys, and employee workforce passwords.
- Manage, monitor, and audit every privileged session—no matter how ephemeral.
- Effectively manage and reduce the entire identity attack surface, spanning Microsoft, Okta, Ping, Salesforce, GitHub, and other domains.
- Intelligently detect and neutralize identity attacks with velocity and precision.
- Satisfy rigorous compliance and forensic requirements by providing easy-to-access reporting on all privileged activity.
- Qualify for cyber insurance by addressing key security controls demanded by cyber insurance providers and policy underwriters.

With BeyondTrust, organizations not only benefit from powerful preventative security layers against both external and insider threats, but also cutting-edge detection and response capabilities to stop in-progress attacks. Learn more.

# Conclusion

This year's report unified two things: it analyzes Microsoft's published vulnerability trends, and it pulls forward the "ghost in the machine" perspective to reflect how modern environments are actually breached, through identity pathways that increasingly include machines and autonomous NHIs. We would encourage readers to unify these perspectives in their own organizations to avoid siloing vulnerability management and identity security, and instead take a holistic view of the threat landscape.

Organizations must prioritize remediation, not only based on the vulnerabilities most likely to enable remote code execution, privilege escalation, and lateral movement, but also by pairing patches with controls that reduce blast radius, including least privilege, governance over secrets and tokens, and robust monitoring over identities.

Until next year, stay secure, keep patched, and protect all privilege pathways.

# Methodology

Every Tuesday, Microsoft issues "Patch Tuesday" and releases security bulletins announcing fixes for any vulnerabilities affecting Microsoft products. The annual Microsoft Vulnerabilities Report (developed and published by BeyondTrust) compiles these releases into a year-long overview and analyzes the data, creating a holistic view of trends related to vulnerabilities.

Until November 2020, Microsoft had been using their own method of sharing CVE details via their Security Update Guide. The former reporting format featured an executive summary for each reported vulnerability that would include the following verbiage:

- Customers / users whose accounts are configured to have fewer user rights on the system could be less impacted than users who operate with administrative user rights.
- If the current user is logged on with administrative user rights, an attacker could take control of an affected system.

From this summary, security researchers could deduce whether any given vulnerability (specifically, the critical ones) could have been mitigated had admin rights been removed from the user.

In 2021, Microsoft shifted methodologies and moved to the Common Vulnerability Scoring System (CVSS). And in 2023, Microsoft continued to use CVSS 3.1 scoring for their vulnerabilities, but began ranking severities based on Microsoft's own Security Update Severity Rating System.

The CVSS methodology allows Microsoft's vulnerabilities to be cross-referenced more easily with third-party bugs, simplifying some analysis, and Microsoft's Security Update Severity Rating System allows each vulnerability to be rated according to the worst theoretical outcome, should that vulnerability be exploited.

One thing that was lost in this new methodology, was consistently listing removal of admin rights as a mitigation for vulnerability risk. While the risks of excess privilege remain very much intact, they are no longer directly highlighted in the data.

Thus, while the statistics on the number of vulnerabilities mitigated by removing admin rights may be absent from this year's report, it's imperative that organizations don't get complacent. Removal of admin rights remains a key piece of a least privilege strategy, as well as for enabling zero trust.

## Accuracy of Vulnerability Data

A number of generalizations have been made for each vulnerability, as follows:

- Each vulnerability was classified with the highest severity rating of all instances of that vulnerability where it appeared multiple times.
- Each vulnerability was classified with the most prevalent type for all instances of that vulnerability.
- Product versions were not taken into account.
- Product combinations were not taken into account.
- Vulnerabilities were counted for both the software and version where appropriate (for example, a vulnerability for Microsoft Edge on Windows 10 is taken as a vulnerability for both Microsoft Edge and Windows).

# Additional Resources

| | |
|---|---|
| **AWARD-WINNING TOOL: Identity Security Risk Assessment** | See the True Privilege™ of every human and non-human identity (including AI agents), and gain a thorough understanding of your identity security posture, along with actionable mitigation steps. |
| **WHITEPAPER: Buyer's Guide for Complete Privileged Access Management (PAM)** | Understand the must-have PAM and privilege-centric capabilities needed to secure identities and access for all identity types (including AI agents). Includes free vendor comparison checklists. |
| **BLOG: AI Agent Identity Governance: Why Least Privilege is the Non-Negotiable Security Control** | Discover why identity governance for AI agents is important for today's organizations, and why it's crucial to enforce least privilege to prevent unauthorized execution on local endpoints. |
| **RESEARCH: AI Agents & Identity Security: Risks, Controls, and Insights** | Explore recent research on how companies are leveraging agentic AI, including adoption trends, concerning risks, and current security strategies and controls. |
| **CHECKLIST: Need a VPN Alternative?** | Take the Remote Access Test to learn if your team has the appropriate secure remote access tools in place to handle a large volume of remote users. |

BeyondTrust

BeyondTrust is the global privilege-centric identity security leader protecting Paths to Privilege™. We are leading the charge in transforming identity security and are trusted by 20,000 customers, including 75 of the Fortune 100, and our global ecosystem of partners.

**Learn more at www.beyondtrust.com**