FIELD EFFECT

2026 Cyber Threat Outlook

Cybersecurity insights from real threats and frontline telemetry

Table of contents

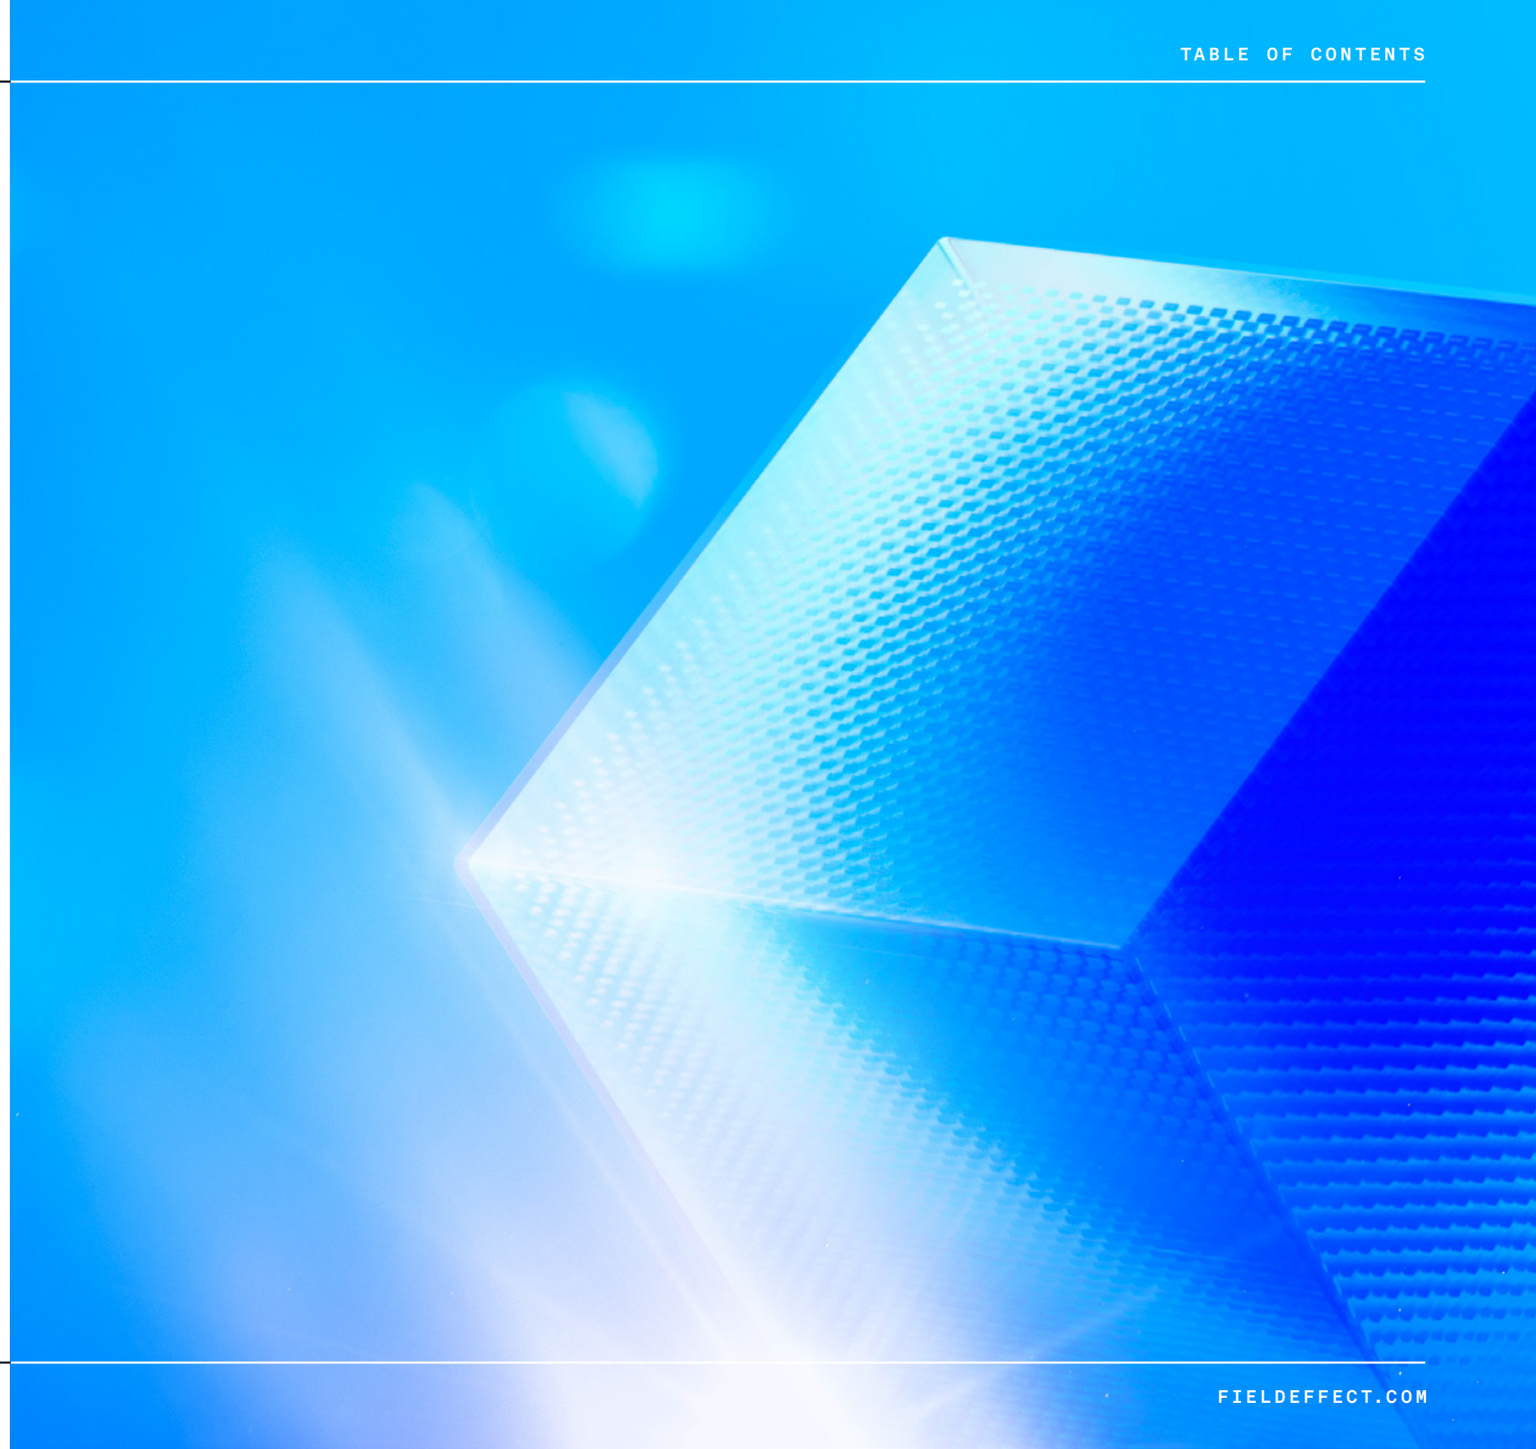

Executive summary

The 2025 cyber threat landscape was shaped by multiple forces that accelerated the speed, scale, and impact of attacks.

Understanding these shifts requires a model that explains both what attackers did, and why their operations intensified.

In cyber threat intelligence, adversaries are best understood through three core factors: their capability to act, their intent to pursue specific objectives, and the opportunity created by environmental or organizational weaknesses.

When these factors align, threats escalate rapidly and with far greater impact, increasing overall threat levels.

In 2025, all three of these elements expanded materially:

- Capability surged as adversaries operationalized AI to reduce the cost, time, and expertise to conduct phishing, develop malware, conduct reconnaissance, and test exploits. Advanced language models let novice actors produce credible phishing content, while sophisticated groups industrialized operations through APIs, automated "support desks," and signed binaries that blended into enterprise environments.
- Intent remained high, fueled by geopolitical tensions, sanctions, trade disputes, and regional conflicts. The result was a steady cadence of espionage, influence operations, disruptive attacks, and opportunistic campaigns. In many cases, these were tied to political events, public uncertainty, and moments of heightened global attention.
- Lastly, opportunity widened rapidly. The swift adoption of AI-enabled browsers, cloud services, autonomous agents, and non-human identities (API keys, tokens, and service accounts) broadened the attack surface without adequate safeguards. Unmanaged browsers, legacy edge appliances, and shadow cloud assets added risk, enabling identity compromise, data leakage, and insider misuse.

Across Field Effect's investigations, reducing opportunity consistently delivered the most immediate and measurable risk reduction. Strengthening identity pathways, securing edge infrastructure, and improving patching practices proved far more effective than attempting to limit adversary capability or intent, which remain largely outside organizational control.

This report is meant to provide a clear, forward-looking analysis of the most significant threats our experts saw in 2025. Backed by Field Effect's frontline intelligence and incident response insights, we highlight critical trends, examine real-world threat campaigns, and offer clear, actionable guidance to help you strengthen defenses for the year ahead.

Primary threat drivers

In 2025, the cyber threat landscape was shaped by a geopolitically driven convergence of state, criminal, and hacktivist operations, and a rapid expansion in the use of generative AI.

Blurring lines across cyber threat actors

Geopolitical friction continued shaping the pace and focus of cyber activity across every region. Conflicts, economic instability, contested elections, sanctions, and trade disputes created high levels of uncertainty that threat actors quickly exploited.

State-aligned groups intensified espionage, influence campaigns, and disruptive attacks, while politically motivated hacktivists fueled instability with waves of DDoS activity, defacements, and data leaks. As governments redirected intelligence and law enforcement resources toward global crises, financially motivated actors faced fewer constraints and escalated their use of ransomware as a primary revenue source.

The year's escalation began when shifting US tariffs created widespread uncertainty, which cybercriminals jumped at the chance to exploit through targeted phishing and impersonation schemes. In Canada, the federal election cycle drove an increase in politically themed phishing, including government impersonation, credential-harvesting attacks against campaign staff, and probing of political party infrastructure.

China's cyber threat in 2025 was defined by a broad, state-aligned ecosystem that blended government agencies, private contractors, and commercial technology firms to support espionage, access operations, and strategic influence.

Leaks tied to Knownsec exposed how contractor networks contributed to large-scale reconnaissance and exploitation efforts, while overall activity demonstrated a continued focus on targeting foreign governments and critical industries through stealthy, long-term intrusion campaigns. China-based threat actors also rapidly weaponized newly disclosed vulnerabilities to gain footholds in enterprise environments before patches were widely deployed. At the same time, geopolitical scrutiny of Chinese technology suppliers intensified, exemplified by Canada's order for Hikvision to cease operations due to national security concerns.

Iran's cyber activity in 2025 showed a pattern of steady, opportunistic, and increasingly adaptive operations shaped by regional tensions and the use of proxy groups. Rather than relying on a single tactic or toolset, Iranian threat actors combined evolving custom malware with rapid exploitation of widely used enterprise software.

Campaigns linked to threat groups like MuddyWater highlighted this shift. New tooling improved evasion, even as broader infection chains continued to depend on familiar techniques such as malicious documents and command-line execution.

Russian cyber threats demonstrated a highly resilient ecosystem that continues blending state-aligned intelligence units, criminal groups, and infrastructure providers operating with a high degree of government tolerance. All year, Russia-based threat actors conducted espionage and access operations against Western governments, defense contractors, and organizations supporting Ukraine, while continuing to exploit zero-day vulnerabilities.

Sustained activity of Russia-based ransomware groups, including families such as Akira, LockBit, Black Basta, continued to target critical industries and government networks worldwide.

North Korea's cyber threat in 2025 was driven by state-directed revenue generation, covert workforce infiltration, and a growing reliance on more interactive and deceptive intrusion techniques. The regime expanded its global IT worker schemes as operatives, often armed with forged identities, infiltrated foreign companies to channel earnings and privileged access back to the capital of North Korea.

Hacktivism shifted toward real-world disruption. Activists moved beyond website defacements to directly manipulate internet-exposed industrial control systems across Canada, including altering water-facility pressure values, tampering with tank gauges, and changing agricultural silo conditions. Similar patterns emerged in the US.

Looking back, 2025 marked another year in which geopolitical tension and cyber activity are inseparable. Strategic, financial, and political motives frequently overlapped, producing a threat landscape where state operations, cybercrime, and hacktivism increasingly converged.

Generative AI amplified adversarial tradecraft

The most significant shift in adversary capability was the increased operational use of generative AI. While AI didn't necessarily introduce brand new attack vectors, it amplified nearly every existing one.

Generative AI models enabled the rapid production of credible phishing content, clean malware code, multilingual lures, and automated reconnaissance. This both boosted the capabilities of advanced actors and lowered the barrier for novice ones.

2025 saw the rise of prompt injection attacks, with malicious instructions embedded in ordinary text redirecting AI-enabled systems to leak data, bypass controls, or execute harmful actions. Meanwhile, ransomware operators adopted faster, more professionalized techniques focused on data theft, and state-aligned actors leveraged AI for exploit development and influence operations.

Cybercriminal marketplaces offered AI-enhanced services such as phishing kits and malware-as-a-service platforms, while attackers exploited public interest in AI by disguising malware as fake AI tools.

Researchers also demonstrated how AI chatbots could be tricked into bypassing safety filters through malicious prompts hidden in fictional narratives, underscoring the need to treat AI systems as part of the attack surface.

As highlighted by Matt Holland, Field Effect's Founder and CEO, AI is also enabling threat actors to treat vulnerability exploitation as a fully automated pipeline. Tasks that once required human effort, such as testing proof-of-concept exploits, identifying misconfigurations, or chaining vulnerabilities, can now be performed programmatically.

The result is a shrinking window between a vulnerability being disclosed and actively weaponized.

Emerging threats in 2025

In 2025, identity was the dominant attack surface, with the edge being the entry, and trust being what adversaries ultimately abuse.

Identity compromise as the primary attack vector

Organizations faced a sharp rise in attacks targeting both human and non-human identities (such as service accounts, API keys, OAuth tokens, and other high-privilege credentials) that often operate with limited oversight.

In September 2025, Field Effect began tracking a Microsoft Teams vishing campaign using Quick Assist to deliver a PowerShell-based web-socket remote access trojan (RAT). Threat actors impersonated internal IT staff through rapidly created onmicrosoft[.]com tenants and rotating "Help Desk" accounts, convincing users to grant Quick Assist access that enabled privilege enumeration and multi-stage malware execution.

Earlier in the year, Field Effect observed similar patterns in a BlueNoroff (Lazarus Group) operation, where state actors impersonated trusted business contacts during scheduled Zoom meetings and used spoofed Zoom-themed domains to trick victims into running malicious "audio repair" scripts. This activity relied on the same principles: exploiting trusted platforms, embedding malicious actions within normal workflows, and using legitimate-looking scripts and update mechanisms to deploy multi-stage malware, harvest credentials, and maintain persistence.

Both campaigns aligned with methods seen across the broader threat landscape. The use of rapidly created M365 tenants, IT-themed personas, and voice- or meeting-based social engineering aligned with the likes of Scattered Spider, Midnight Blizzard, LAPSU$-style groups, and multiple ransomware operators. A reliance on PowerShell stagers, AMSI bypasses, RMM tools, and credential harvesting echoed techniques used by Akira and in other financially motivated intrusions.

These cases reflect a defining trend: threat actors across multiple clusters are increasingly weaponizing identity, trust, and collaboration tools to gain initial access. Identity has become the primary attack surface, and trusted enterprise platforms, such as Microsoft Teams, Zoom, Quick Assist, RMM utilities, and Microsoft 365 cloud identities, are being exploited to bypass defenses, establish persistence, and escalate access.

Generative AI further amplified this shift by enabling highly convincing phishing, vishing, and social engineering campaigns that fuel identity-based ransomware and credential theft. Attackers increasingly targeted identity providers, exploited weaknesses in SSO and MFA, and used infostealer malware and recirculated breach data to bypass authentication entirely.

Several systemic drivers consistently weakened authentication integrity and expanded compromise opportunities, including password reuse, unmanaged devices, AI-generated lures, and a growing number of non-human identities.

Credentials, tokens, SSO abuse, MFA fatigue, session hijacking, and privilege escalation have become the primary mechanisms through which attackers gain meaningful access—far outpacing traditional network exploits.

Software supply chain attacks also increasingly centered on developer identity compromise. Threat actors used phishing and credential theft to access trusted maintainer accounts. This was seen in the Shai Hulud campaign, where attackers compromised maintainers, stole NPM and GitHub tokens, republished malicious packages, moved laterally across repositories, and automated propagation using legitimate tools.

This shift underscores a broader 2025 trend: protecting the identity of both human and machine has become the foundation of modern security.

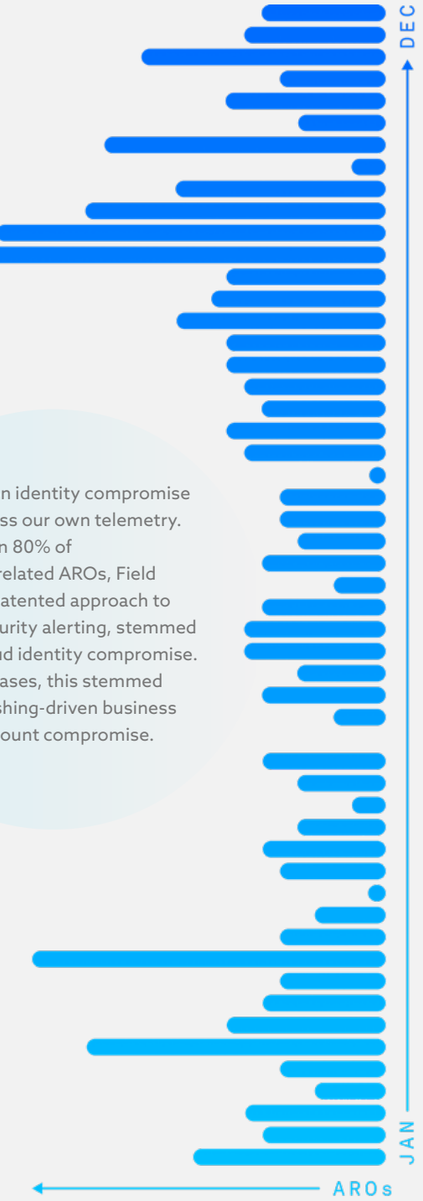

A focus on identity compromise held across our own telemetry. More than 80% of incident-related AROs, Field Effect's patented approach to cybersecurity alerting, stemmed from cloud identity compromise. In most cases, this stemmed from phishing-driven business email account compromise.

The disappearing perimeter and the rise of edge-focused attacks

Throughout 2025, threat actors increasingly targeted the edges, exploiting a space where visibility was limited, patching lagged, and the fine line between “inside” and “outside” no longer held.

Routers, VPNs, firewalls, cloud-exposed services, and other perimeter-adjacent systems became high-value entry points. This was especially true as critical vulnerabilities in edge devices, DNS services, and major web browsers saw rapid proof-of-concept releases and, in some cases, exploitation before organizations could respond.

These internet-facing systems, often managed by small IT teams or third-party providers, expanded the attack surface alongside persistent software-supply-chain risks.

Dependencies, libraries, and managed services introduced security gaps that were difficult to track, pushing organizations toward stronger asset discovery and better dependency management. The perimeter no longer has a clear boundary and is becoming increasingly difficult to defend.

For many smaller organizations, this reality materialized during the SonicWall exploitation campaign, which exposed just how easily attackers could bypass weakened perimeter defenses.

The key takeaway here is how quickly one weakness at the edge can unravel broader defenses, enabling stealthy access, identity compromise, and rapid escalation. Going forward, treating edge devices as high-value assets, subject to the same rigor as domain controllers, should become standard operating practice.

2025’s perfect storm: SonicWall exploitation and the Akira ransomware campaign

In 2025, Field Effect documented a sustained exploitation campaign targeting SonicWall SSL VPN and management interfaces. The attack leveraged valid credentials, many of which were tied to historic SonicWall compromises. The vendor later linked this campaign to a previously disclosed and patched vulnerability, CVE-2024-40766.

Akira ransomware operators used unrotated credentials to authenticate directly into high-privilege systems, including Domain Controllers, without triggering a typical security alert.

After gaining access, the attackers rapidly disabled security controls, conducted targeted and broad file searches, and exfiltrated data using automated WinRAR tooling. They then deployed encryption as part of a double-extortion attack, one where victims are pressured to pay both for decryption and to prevent the stolen data from being publicly disclosed.

Many affected organizations had not applied the patches for CVE-2024-40766, despite SonicWall releasing fixes prior to the surge in attacks. Others had applied the patches but failed to rotate critical credentials that had been exposed while the systems were vulnerable, which were later reused to compromise their networks. Field Effect assesses that it is likely an Initial Access Broker (IAB) initially collected these credentials, which were subsequently leveraged by the Akira group.

It was a perfect storm: sophisticated attackers armed with valid administrative credentials, widespread credential reuse, and unpatched systems.

Exploiting trust across people, platforms, and processes

Throughout 2025, Field Effect investigations revealed a consistent pattern: threat actors advanced their objectives by inserting themselves into the trusted spaces where people, tools, and workflows intersect.

Collaboration platforms such as Teams and Zoom became high-value vectors, with adversaries using fake meeting invites, impersonated tenants, spoofed domains, and voice-based social engineering to deliver malware, deploy remote-access tools, and harvest credentials under the guise of routine communication.

This dynamic appeared in campaigns targeting administrative and support tooling. In one investigation by Field Effect, adversaries abused trust in the remote support platform SimpleHelp, using it as a delivery channel for backdoors. Because SimpleHelp is normally associated with IT assistance, the activity appeared to originate from a trusted support tool and malicious actions blended in with normal behavior.

This same pattern extended beyond support utilities into the tools administrators rely on to find and download software. One campaign, Thunderstruck, demonstrated how threat actors can move even earlier in the workflow, compromising the discovery process itself. By impersonating RVTools (a trusted VMware administration tool) in malicious search engine ads, threat actors redirected administrators to a fake installer that deployed the Thundershell payload, turning a routine search for a trusted utility into an entry point for compromise.

In the Grixba-Play Ransomware campaign, attackers abused trust in security vendors by disguising malicious components as SentinelOne tools. Threat actors gained execution by mimicking the appearance and behavior of a well-known EDR product. In this case, the supposed defender became the attacker's disguise.

Insider threats also increasingly centered on the abuse of trusted access, fueled by remote work, global hiring, and cloud-first operations. Field Effect uncovered numerous cases where individuals used valid credentials and routine workflows to mask malicious activity, with several incidents escalating to attempted data theft or sabotage once scrutiny began.

In all these events, threat actors consistently exploited the trust that organizations place in their people, tools, and processes. The success of these intrusions depended on the attacker's ability to appear legitimate, either by convincing users that a support technician was genuine, that a meeting invite came from a known contact, that a software download was authentic, or that a system alert required urgent action.

One organization discovered a new employee falsely claiming to be in Canada while they were operating from a country in Africa. When confronted, the newly hired remote worker attempted to deploy USB-based malware but was blocked by Field Effect MDR before any true damage was done. Post-incident reviews found other staff at the company also misrepresenting their locations.

Similar patterns appeared in widely publicized cases in which companies unknowingly hired North Korean operators, revealing a broader surge in applicants using fabricated identities or rotating personas to obtain trusted access.

The threats that stuck around

Ransomware remained a steady, long-running threat last year, reinforced by the same factors that have driven its success for years: it's profitable, scalable, and difficult to fully contain. Meanwhile, human and technical vulnerabilities continued to provide the easiest entry points for attackers.

Ransomware continued: Factors evolving the threat in 2025

Despite more than a decade of attention, ransomware remained one of the most reliable and scalable tools available to cybercriminals. With much of the ecosystem now automated and profit potential high, ransomware still offers a strong return on investment.

The continued exploitation of VPN appliances and other edge-exposed systems remained a core enabler of ransomware operations. High-profile incidents such as the SafePay ransomware compromise of Ingram Micro, enabled through a GlobalProtect VPN intrusion, demonstrated how edge exploitation quickly leads to identity abuse and disruption.

As ransomware groups looked for higher-impact opportunities, virtualization platforms became a more regular and intentional target in 2025. Groups such as Akira expanded their tooling to encrypt workloads running on Nutanix AHV and other highly privileged infrastructure.

Ransomware operators also benefited from the accelerating pace of vulnerability exploitation, exploiting critical flaws in platforms such as Oracle E-Business Suite and SonicWall appliances within days of disclosure.

Finally, the expansion of ransomware-as-a-service ecosystems, accelerated by increasing automation and AI-driven tooling, further lowered barriers to entry and increased the speed and scale of operations. New families such as Warlock showed how quickly threat actors can adopt shared tooling and affiliate models, ensuring ransomware remains a persistent threat across industries.

These examples made clear that ransomware continues to evolve in ways that exploit systemic weaknesses in identity, trust, and infrastructure.

Human decisions: The real drivers of cyberattacks

Even as technical defenses continued to advance, adversaries relied heavily on human error, misplaced trust, and operational shortcuts to gain initial access and escalate their impact.

Social engineering remained a reliable entry point for intrusions, driven by credential theft, password reuse, and multifactor authentication (MFA) fatigue across all organizations. What changed was how these tactics were executed.

Rather than depend on malware-laced attachments or obvious phishing links, adversaries posed as IT support, mimicked familiar workflows, and delivered support-style instructions intended to convince users to run commands, download tools, or grant access directly.

Field Effect observed this shift most clearly in campaigns like ClickFix, first tracked in early 2025. In these cases, adversaries impersonated internal IT teams and used fake CAPTCHA prompts to guide users into manually executing malicious PowerShell commands.

These commands deployed remote access tools and infostealers such as AsyncRAT, establishing persistence and enabling follow-on activity without triggering conventional exploit-based defenses. This marked a broader strategic shift: when adversaries couldn't bypass security controls, they simply asked users to do it for them.

Several other incidents in 2025 highlighted how unintentional actions can create serious security gaps. Routine oversights, such as failing to apply patches, misconfiguring access controls, or downloading tools from unverified sources, frequently opened the door for threat actors long before any sophisticated technique was deployed.

But, at the same time, people remain one of the most effective lines of defense. Field Effect observed this firsthand when an HR team member identified a subtle discrepancy in a candidate's video background during a routine interview. Pausing the conversation and escalating the concern internally prevented what may have been an attempt to gain unauthorized access through the hiring process.

While identity and integration controls are certainly critical, human vigilance remains equally essential in detecting subtle anomalies that automated systems might miss.

Vulnerabilities remained a consistent focus

Throughout 2025, threat actors continued to rely on outdated software, firmware, and exposed edge infrastructure as dependable points of entry. There's no doubt that this is a pattern that will follow well into 2026 and beyond.

But this pattern only captures a fraction of the real risk landscape. The technologies creating the most exposure are actually often ignored: legacy systems, neglected tools, and products from vendors that never issue patches.

Industry data challenges make the picture even murkier. A common pattern in our security intelligence work is that after a vendor is acquired, vulnerability data starts appearing under the new name while existing deployments inside organizations remain listed under the old one. The same product ends up split across multiple identifiers, and because CPE matching relies on exact strings, those legacy entries are often missed by IT departments. It's a structural gap that leaves real exposures hidden.

This gap has widened with the growth of shadow IT. Without clear policies and oversight, employees adopt AI tools, browser extensions, niche browsers, and specialized apps that sit entirely outside formal security controls. These tools introduce risks that never appear in public datasets and often evade traditional monitoring.

All of this creates a distorted sense of where risk actually lives. Vendors with mature security programs tend to appear more often in advisories simply because they disclose more. Meanwhile, unpatched tools, abandoned components, shadow IT, and products with mismatched identifiers remain largely invisible.

Field Effect MDR telemetry showed a steady, ongoing presence of vulnerability-related risk across customer environments. The monthly spikes* in the graph below reflect our careful tracking of events like Patch Tuesday and other regular vendor update releases.

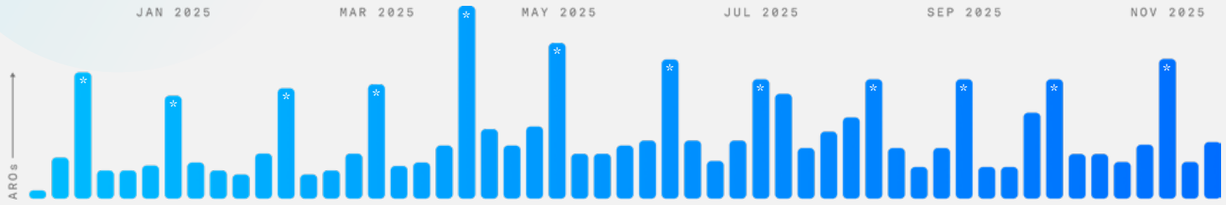

THIS YEAR'S VULNERABILITY TRENDS:

Most exploited attack surface categories

01 EDGE-DEVICE AND NETWORK APPLIANCES
Threat actors continued to prioritize internet-facing appliances, including WatchGuard, SonicWall appliances, Cisco, Citrix appliances, TP-Link, ASUS WRT and other end-of-life routers, DrayTek devices, Axis and Hikvision surveillance systems, and Red Lion Sixnet.

02 HIGH-PRIVILEGE ENTERPRISE MANAGEMENT SYSTEMS
Due to the administrative access they provide, high-privilege infrastructure platforms remained valued targets. These include CentreStack and Triofox, ConnectWise and N-able RMMs, Commvault, Fortinet, Paessler, Bomgar and SimpleHelp remote-access tools, and Palo Alto GlobalProtect.

03 MICROSOFT, BROWSER, AND OS-LEVEL ZERO DAYS
Microsoft and browser ecosystems were repeatedly targeted due to their ubiquity and privileged access. Some exploited vulnerabilities we saw affected Microsoft SharePoint, Chromium and Firefox, FreeType, VMware Tools, and Aria Operations.

04 AUTHENTICATION AND SSO BYPASS VULNERABILITIES
Identity-related vulnerabilities, such as FortiCloud SSO bypasses across FortiOS, FortiWeb, and FortiProxy, remained among the most damaging due to their ability to grant administrative access without malware.

05 DEVELOPER, CMS, AND AUTOMATION ECOSYSTEM
Threat actors increasingly targeted developer tools, CMS platforms, and workflow automation systems such as Craft CMS, Kentico Xperience CMS, XWiki, Adminer, Apache Tomcat, Adobe AEM, and Adobe ColdFusion.

Hardening defenses for the future

Strong cyber resilience depends on steady execution of the fundamentals, not chasing the newest tools or trends. Even as AI capabilities advance, most defense failures still stem from the same long-standing weaknesses.

When the basics aren't consistently enforced, even well-funded security programs struggle to deliver meaningful protection. The path forward is grounded in embedding best practices into daily workflows, ensuring controls are consistently applied, and scaling these fundamentals across the entire environment.

Stay ahead of AI-driven threats

Clear guardrails around which AI tools are approved, how sensitive data can be used, and training users to recognize AI-generated material lowers the risk of data leakage, fraud, and social engineering.

Beyond policy and awareness, organizations benefit from behavior-based detection that can keep pace with AI-accelerated attacks. Field Effect MDR provides continuous monitoring across endpoints, networks, cloud environments, and identities, surfacing early indicators such as privilege escalation, anomalous logins, or API misuse.

Improve monitoring and detection

As threat actors increasingly rely on legitimate tools and valid credentials, effective detection requires continuous, cross-environment visibility and behavioral analysis—not isolated logs or static alerts.

Field Effect streamlines protection across endpoints, networks, and cloud services with real-time monitoring and full visibility into user and organizational behavior. It surfaces the most critical information through clear, actionable ARO alerts that explain what happened, why it matters, and how to respond.

Field Effect rapidly detects and responds to suspicious activity early in the attack lifecycle, validating and containing threats on an organization's behalf while providing expert guidance on next steps. This combination of real-time monitoring and hands-on support is essential in today's evolving threat landscape.

Reduce exposure at the edge

With internet-facing infrastructure becoming some of the most heavily targeted ground in 2026, organizations need clear ownership and fast response cycles. Field Effect MDR supports this by surfacing outdated systems, unused software, and misconfigurations that increase exposure, with AROs providing early visibility and practical, step-by-step remediation guidance.

When new vulnerabilities appear, our team issues timely alerts and instructions, including temporary safeguards when patches aren't yet available.

Harden identity controls and monitor privileged access

Privileged accounts remain especially attractive for attackers seeking lateral movement and escalation. Enforcing least privilege, monitoring administrative activity, and adopting just-in-time access models help limit the blast radius when compromise occurs.

Strengthening authentication through universal MFA, password managers, stronger identity verification, and preparing for passwordless authentication helps reduce exposure to these threats.

Field Effect MDR reinforces these practices by continuously monitoring authentication activity, privilege changes, and unusual login patterns across cloud and on-prem environments.

Protect data throughout its lifecycle

Regardless of how an adversary gains access, data theft remains a primary objective. Strong encryption, strict access controls, thoughtful retention policies, and minimizing unnecessary data collection all help ensure that even if a breach occurs, the potential impact is contained.

Field Effect MDR fortifies these practices by monitoring for early signs of data staging, unusual access patterns, and suspicious file activity across endpoints, networks, and cloud environments.

Prepare for fast-moving ransomware

Ransomware operations will accelerate as AI improves reconnaissance and initial access. Field Effect MDR is designed to catch ransomware precursors, such as credential harvesting, lateral movement, and suspicious privilege escalation, to block malicious activity before encryption has the opportunity to begin. When early indicators surface, we provide clear guidance to contain the threat and prevent escalation.

What's next in 2026 and beyond

In 2026, the cyber threat landscape is likely to intensify as geopolitical instability continues to erode the boundaries between state operations, cybercrime, and hacktivism. Advances in generative AI will likely accelerate this convergence, enabling adversaries to weaponize vulnerabilities and craft highly convincing social-engineering campaigns at greater speed and scale. As AI-powered automation becomes more widespread, threat actors will continue industrializing reconnaissance, exploit development, and social engineering.

Identity will likely remain the main battleground. Credential-driven attacks are expected to become more sophisticated as adversaries use AI to convincingly mimic employees, IT staff, and automated systems. Many organizations may find themselves overwhelmed by the rapid growth of non-human identities that now outnumber human users and create an unclear attack surface.

At the same time, the edge is likely to remain one of the most contested spaces in cybersecurity. As more services are pushed to the perimeter, gaps caused by unclear responsibility for maintaining or securing these systems, along with slow patching and limited visibility, will likely further separate what's exposed from what's protected.

Across all of this, threat actors are likely to continue exploiting trust more effectively than software flaws. Impersonation, poisoned search results, malicious extensions, and spoofed support interactions may become even more central tactics. The line between insider and adversary is likely to blur further as covert operators take advantage of good-intentioned business processes.

Ransomware is here to stay, and will likely accelerate. The same dynamics that sustained high activity in 2025, such as state tolerance, mature criminal ecosystems, low barriers to entry, and profitability, are a long-term trend. The overlap between financially motivated groups and state-aligned actors is likely to deepen, reinforcing ransomware's dual role as both a profit mechanism and a geopolitical tool.

The human element will always be a liability and critical line of defense. Attackers will continue exploiting people through workflow-based deception, fake support interactions, and subtle manipulations that convince users to bypass their own controls. At the same time, organizations that invest in strong security culture, clear escalation paths, and modern training are likely to gain a meaningful advantage.

By the end of 2026, vulnerability management is due to shift from a routine, scheduled task to something more continuous and urgent. Threat actors are expected to focus even more on overlooked systems, authentication bypasses, and high-privilege infrastructure. At the same time, shadow IT such as unapproved AI tools, niche browsers, and extensions adopted by teams without security oversight will add new exposures that traditional processes are not designed to catch.

In a threat landscape shaped by automation, AI-driven attacks, and fading trust boundaries, being able to defend at close to machine speed, while still relying on human judgment, may become the clearest sign of real cyber resilience.

A note from the authors

This report was developed by several of Field Effect's security experts, including leaders across our Security Services, Threat and Risk Intelligence, and Professional Services teams, with contributions from experts across the globe. It reflects the work of cybersecurity investigators, analysts, and researchers who work every day to keep businesses secure.

Our goal was to provide a clear, practical view of how the threat landscape evolved and what that means for the year ahead. Every insight in this report is rooted in real-world incidents, validated telemetry, and the threats we observed, investigated, and responded to firsthand.

On that note, threats don't pause between reporting cycles and neither do we. Our teams continue to track evolving attacker behavior, uncover new campaigns, and share what matters most in the moment. You can stay current with our weekly threat roundup, delivered straight to your inbox. Sign up to stay informed all year long.

FIELDEFFECT.COM/NEWSLETTER

FIELD EFFECT / MDR

Complexity out. Clarity in.

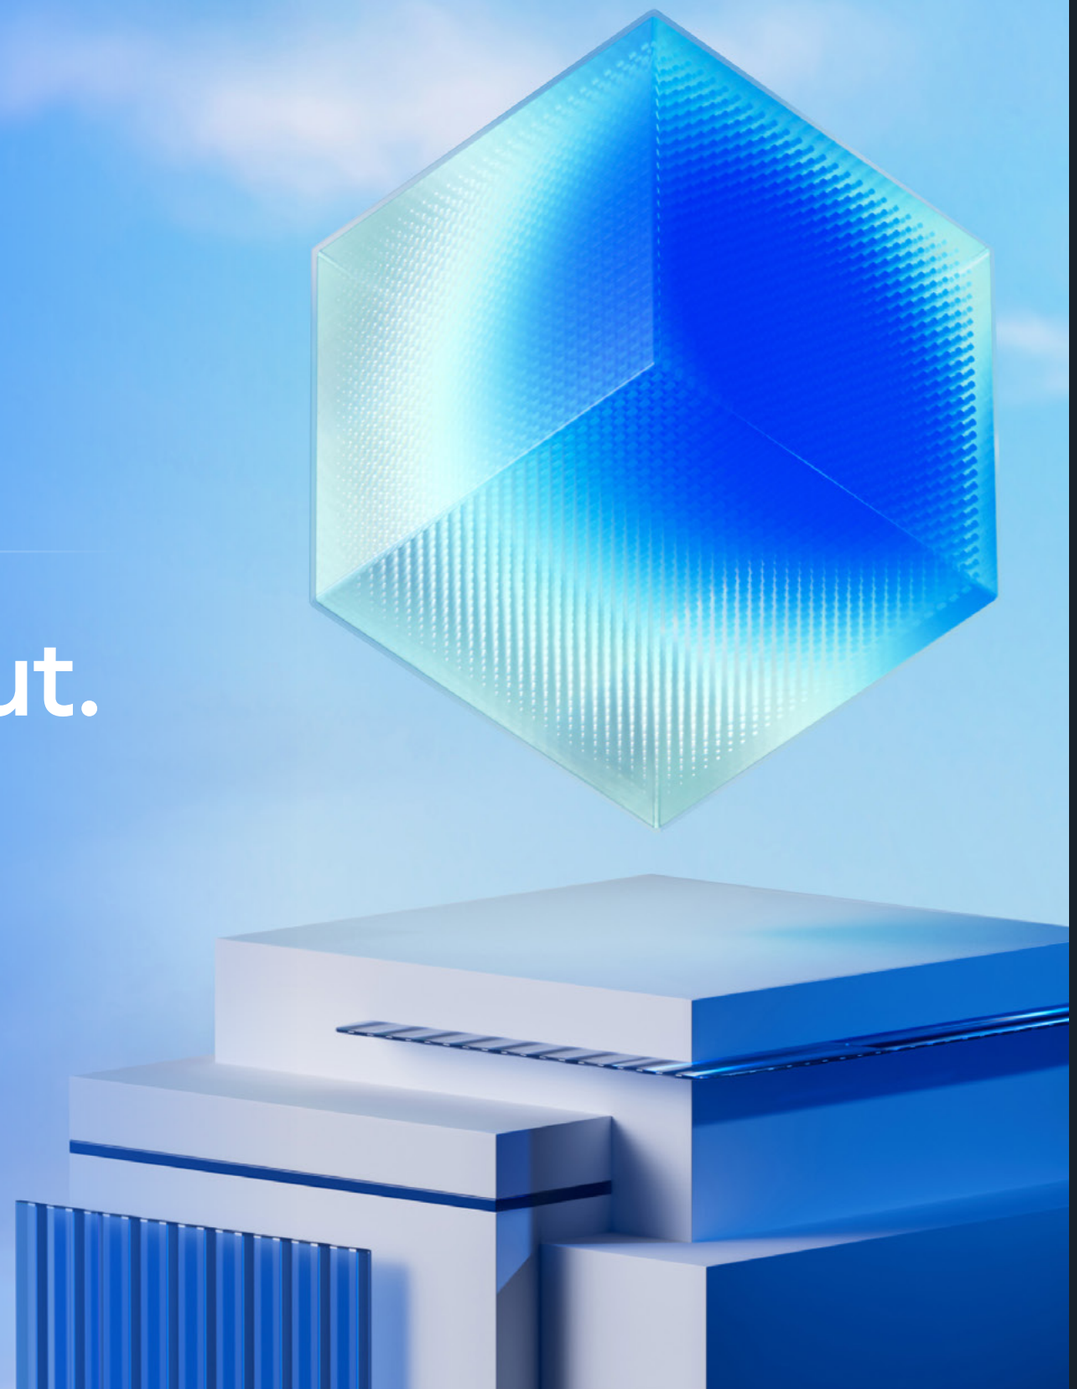

About Field Effect

Every business deserves powerful protection from cyber threats.

Field Effect's cybersecurity solutions were purpose-built to prevent, detect and respond to threats for clients of all sizes. We take on the complexity behind the scenes and deliver a solution that's sophisticated where it matters, and simple everywhere else. Consolidate your tech and eliminate the noise while empowering users of all technical backgrounds to confidently navigate cybersecurity and avoid disruptions.

Contact our team today.

Email:
letschat@fieldeffect.com

Phone:
+1 (800) 299-8986

FIELDEFFECT.COM