TrendAI™ Research

TrendAI™ 2026 Cyber Risk Report

Table of Contents

Building on the findings of the 2025 Cyber Risk Report, TrendAI™ Research continues to track risk insights across enterprise environments globally to better equip organizations with the data they need to stay ahead of an evolving threat landscape.

To achieve this proactive approach, we draw on data from the TrendAI Vision One™ Cyber Risk Exposure Management (CREM) solution, which protects organizations' digital assets by evaluating risks across the attack surface, prioritizing them, and implementing appropriate countermeasures.

The Cyber Risk Index (CRI) sits at the center of CREM, quantifying an organization's overall security risk by consolidating individual asset and risk factor scores. Our research has found that organizations with a CRI above the average are more likely to suffer attacks than those with a lower score. CREM calculates this index by multiplying each asset's attack exposure and security configuration against its criticality, producing a score between 0 and 100 across three risk levels:

- Low risk (0 – 30), where organizations are considered relatively secure and immediate action is generally not required.
- Medium risk (31 – 69), where several risk factors need to be addressed and countermeasures should be implemented.
- High risk (70 – 100), where organizations face severe exposure and prompt, robust security measures are essential.

A risk-based approach remains our guiding priority for improving enterprise security posture. By understanding where exposures concentrate, how attackers exploit them, and how quickly defenders respond, organizations shift from reactive defense to proactive resilience.

1

About our data

This year's report introduces a new dataset drawn from telemetry generated by the TrendAI Vision One™ Attack Path Prediction capability, providing a forward-looking complement to the Cyber Risk Index scores and event detections that anchor our annual analysis. Where the CRI measures the current state of an organization's risk posture, Attack Path Prediction maps the routes an adversary could plausibly follow through that environment, connecting exposed entry points, active risk events, and vulnerable target assets into end-to-end attack sequences. By surfacing these predicted paths before they are exploited, the dataset shifts the analytical frame from exposure measurement to adversary simulation, giving security teams a concrete basis for prioritizing which risks to address first based on how attackers are likely to move.

This transitional report builds on the previous Cyber Risk Report and draws primarily on 2025 telemetry; unless otherwise noted, findings are based on that year. Some sections use the latest available dataset snapshots where full-year data was not yet available; these are intended to provide directional insight rather than a complete annual measurement.

As 2026 telemetry matures, we expect the 2027 report to expand this view with more AI-related risk and threat datasets, building on the baseline risk, detection, vulnerability, and attack-path findings introduced here.

This report draws on telemetry from organizations using the TrendAI Vision One™ platform, together with TrendAI™ threat intelligence. Figures reflect observations across our customer base rather than a statistically representative sample of any industry, region, or the threat landscape overall, and are therefore presented as a directional signal. Explanations of why metrics move or differ reflect our analysts' informed judgment rather than statistically tested or causal findings. These insights are meant to support awareness and risk prioritization, not to benchmark individual organizations or to inform compliance, legal, or investment decisions. Information is current as of publication and provided as is.

2

Risk data

Overall average Cyber Risk Index for 2025

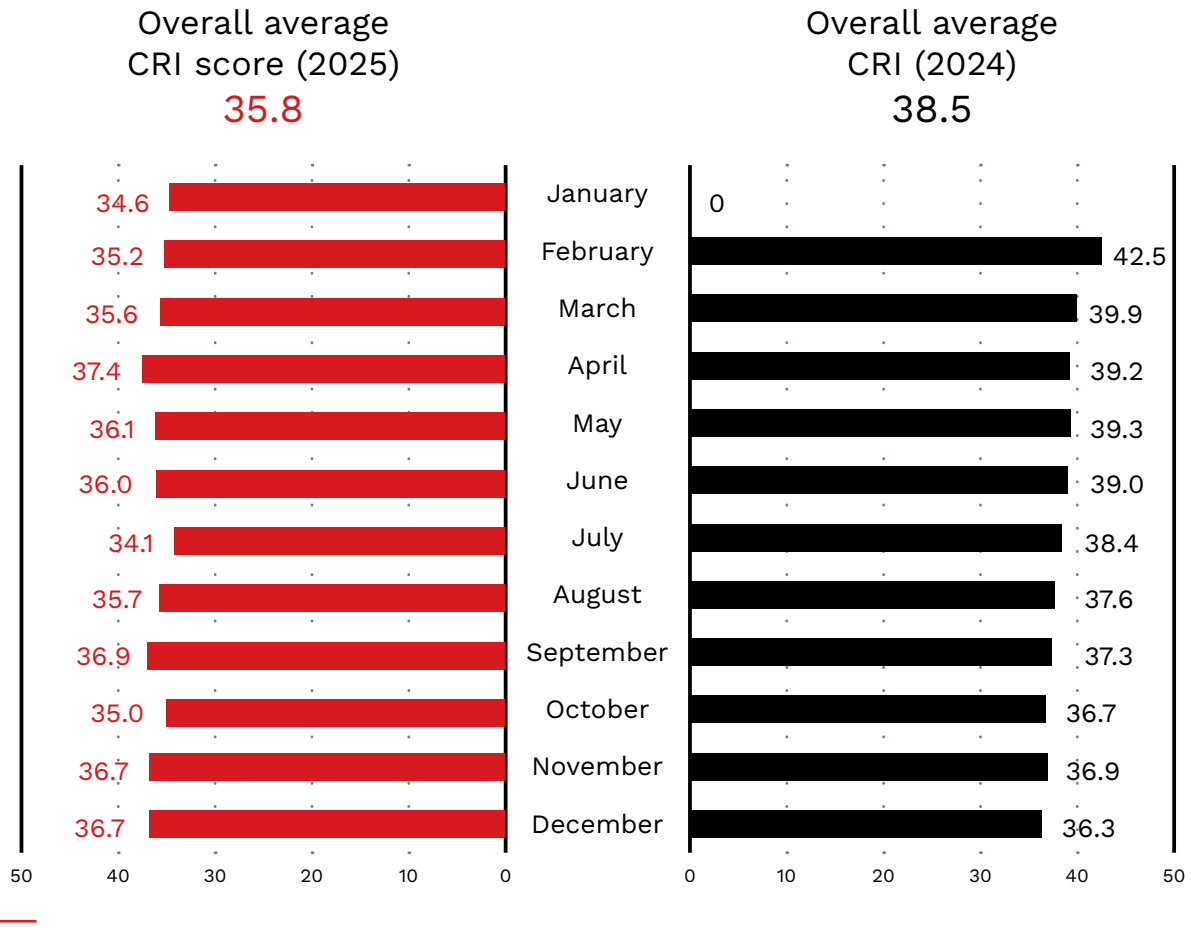

Figure 1. A side-by-side comparison of the Cyber Risk Index overall annual average and monthly breakdown between 2024 and 2025. Figures have been rounded to one decimal place.

(Note: Data for January 2024 is unavailable because the dashboard algorithm was updated at the end of that month with a weight summation method that affects CRI computation.)

The 2025 annual average CRI sits at 35.8, a meaningful improvement from the 2024 average of 38.5 (February to December 2024). This continues the downward (improving) trend seen throughout 2024 and suggests enterprises are making further progress operationalizing cyber risk management, but it should be noted that 35.8 still falls firmly in the medium risk band.

Monthly movement from our telemetry in 2025 shows some volatility compared to 2024's steady decline. The year starts relatively low at 34.6 in January, spikes to a high of 37.4 in April, dips back to a low of 34.1 in July, then climbs again to settle at 36.7 for both November and December.

This oscillating pattern contrasts with 2024's broadly consistent month-on-month improvement from 42.5 down to 36.3, and may suggest that while organizations have improved their baseline posture, they are struggling to sustain continuous

improvement, potentially reflecting seasonal resourcing constraints or reactive rather than proactive risk management cycles.

The April spike is worth investigating specifically. It may correlate with the start of a new fiscal or budget year for many organizations when new projects and cloud deployments introduce fresh risk before controls catch up.

Every organization in our telemetry, regardless of size or sector, still falls within the medium risk band. Progress is real, but the work is far from done.

Top 10 industries with the highest average CRI

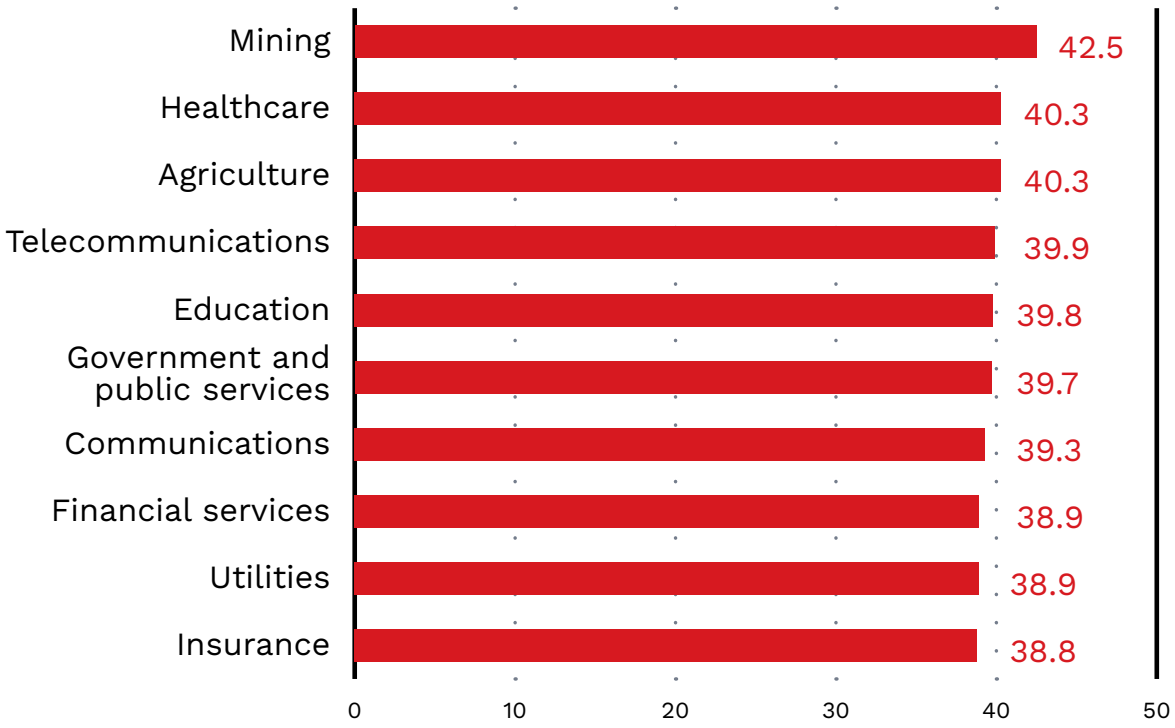

Figure 2. The top 10 industries with the highest annual average CRI for 2025 (January – December). Figures have been rounded to one decimal place; ranking reflects unrounded figures.

The 2025 telemetry shows a reshuffling of the industry rankings from 2024. This section looks at the overall CRI for each of the highest-ranking industries, and possible factors that contributed to their respective scores.

Mining tops the risk rankings for the first time in 2025—a sector absent from last year's top 10 now carries the highest average CRI of any industry.

The mining sector's emergence at the top of the risk rankings likely reflects a combination of accelerating digital transformation and historically underdeveloped security infrastructure. Mining operations are increasingly deploying connected sensors, autonomous vehicles, and remote monitoring systems across geographically dispersed sites—expanding the attack surface rapidly.

At the same time, the sector's top risk event being email-borne threats rather than cloud app access suggests that basic security hygiene, including user education and email filtering, may not have kept pace with operational technology (OT) adoption. Legacy OT systems that were never designed for internet connectivity are now being networked, creating vulnerabilities that are difficult to patch without operational downtime.

Healthcare remains in the top three for the second consecutive year, its risk driven by a complex device estate, slow patch cycles, and persistent identity management gaps.

The healthcare industry's sustained high CRI reflects structural challenges that are difficult to resolve quickly. The sector manages an exceptionally diverse and often legacy device estate (from medical imaging equipment to patient monitoring systems), many of which run outdated software that cannot be easily patched without regulatory approval or risk to patient care continuity.

The 2024 mean time to patch (MTTP) data, which shows healthcare as the slowest-patching industry at 41.5 days, underscores this. Device Control Settings ranks as the industry's top misconfiguration event, consistent with the challenge of governing thousands of connected endpoints. Risky Cloud App Access ranks as its top risk event, which may reflect the rapid adoption of cloud-based tools outpacing the governance frameworks meant to secure them.

Agriculture's risk profile is inseparable from its technology transformation: the same internet of things (IoT) devices and automated machinery driving operational efficiency are expanding the attack surface faster than security controls can follow.

Agriculture's persistently high CRI reflects the sector's ongoing digital transformation, which has outpaced its security maturity. The adoption of precision farming technologies, connected irrigation systems, automated harvesting equipment, and supply chain management platforms has dramatically expanded the attack surface over a short period. Device Control Settings ranks as its top misconfiguration event, consistent with this IoT proliferation.

Two Data Loss Prevention events—email and cloud app violations—appear prominently among its risk events, suggesting that sensitive operational and supply chain data is being transmitted without adequate controls. This risk is compounded by the sector's heavy reliance on third-party vendors and service providers, who may introduce additional vulnerabilities.

Telecommunications enters the top five for the first time in 2025, with access to risky cloud apps and stale identity accounts among its primary risk drivers—challenges that mirror the broader enterprise landscape but carry amplified consequences given the sector's critical infrastructure role.

Telecommunications organizations are undergoing significant infrastructure transformation, migrating toward cloud-native architectures, deploying 5G networks, and managing increasingly software-defined infrastructure. This transition naturally introduces new risk vectors faster than security controls can be established.

Suspicious Connection Service Settings, which ranks as its top misconfiguration event, is particularly notable for a sector whose core business is managing network connections, suggesting that even network-native organizations struggle to maintain optimal security configurations at scale. The combination of large, complex user directories and rapid infrastructure change makes stale account management a persistent challenge.

Education has made the most measurable progress of any sector, dropping from a quarterly peak CRI of 45.1 in early 2024 to 39.8 in 2025. The improvement is real, but with stale accounts and weak authentication policies still topping its risk event list, the sector's identity management challenge remains unresolved.

Education's improvement is likely the result of increased regulatory pressure, greater awareness following high-profile breaches in the sector, and targeted remediation efforts prompted by the 2024 report's findings. However, the factors driving its remaining risk are likely deeply structural. Educational institutions manage exceptionally large and fluid user populations as students, faculty, and staff who cycle in and out regularly, which makes account lifecycle management inherently more complex than in most sectors.

Budget constraints, particularly in public institutions, limit investment in automated identity governance tools. The prevalence of personally owned devices connecting to institutional networks further complicates endpoint control, which might explain why Non-Optimized Security Agent Password Unlock Settings tops its misconfiguration list.

Government and public services organizations carry a CRI of 39.7 in 2025, a marginal improvement from last year's 40.3. This sector should work to lower its CRI at a much steeper curve as consequences of a successful attack extend beyond operational disruption to public safety and national security.

Meanwhile, the communications industry has improved from a 2024 average of 41.6 to 39.3 in 2025, one of the stronger year-over-year improvements in the top 10, though email-borne threats and data loss incidents continue to dominate its risk event profile.

Financial services organizations carry a CRI of 38.9 in 2025, continuing a gradual improvement from 2024. With accounts that have multi-factor authentication (MFA) disabled among its top risk events, even the most regulated and security-conscious sector is not immune to foundational identity hygiene failures.

Utilities enters the top 10 for the first time in 2025—a sector where information technology and operational technology convergence is creating new attack surfaces that traditional security frameworks were not designed to address.

Insurance rounds out the top 10 with a CRI of 38.8, down from 41.0 in 2024. MFA-disabled accounts and Risky Cloud App Access persist as top risk events, and a uniquely diversified multi-cloud environment is introducing governance complexity that may be slowing remediation. Its use of Google Cloud Platform (GCP) alongside the Amazon Web Services (AWS) and Microsoft Azure footprints more common in other sectors adds to that complexity. Knowing the cost of risk and operationalizing its reduction at scale remain two very different challenges.

Average CRI by company size

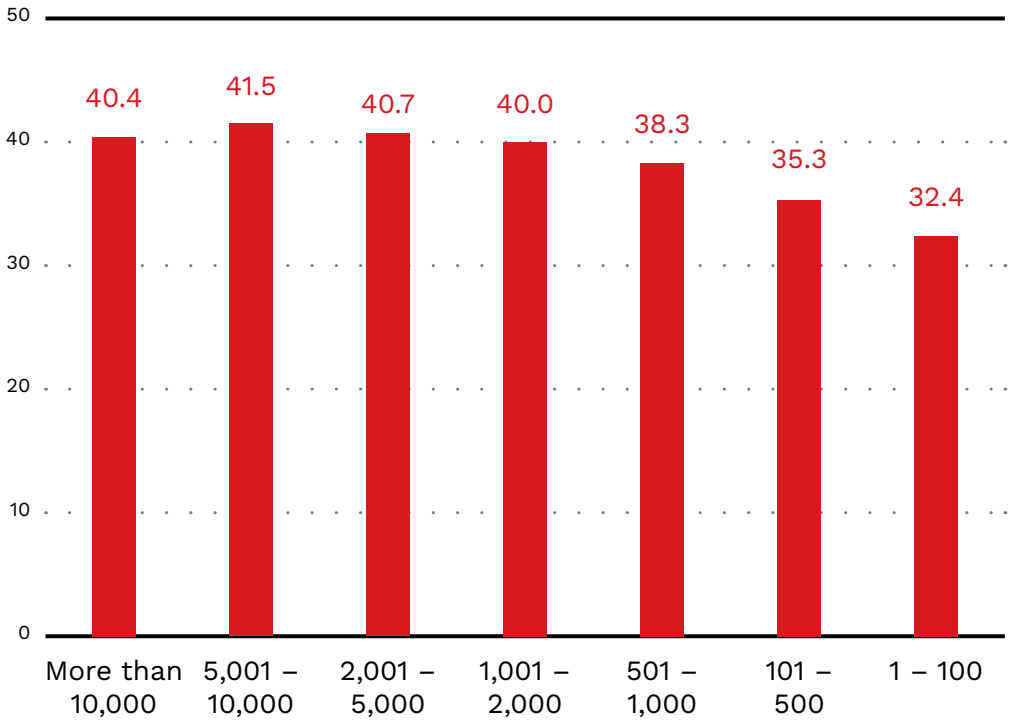

Figure 3. A breakdown of the annual average risk scores for each organization size by employee count for 2025 (January – December). Figures have been rounded to one decimal place.

Small organizations (100 or fewer employees) remain the lowest-risk band overall, but 2025 marks the only size segment to record a CRI increase year-over-year. This uptick likely reflects growing attacker interest in small businesses as entry points into larger supply chains, combined with limited dedicated security resources to respond to an expanding threat landscape. With lean IT teams often managing security as a secondary responsibility, even modest increases in attacker attention can meaningfully increase risk.

Organizations with 501 to 1,000 employees appear to show steady improvement, likely reflecting increased adoption of managed security services and platform-based security tools that bring enterprise-level capabilities within reach. However, this segment sits within the medium risk band at 38.3.

Organizations with 1,001 to 2,000 employees recorded one of the stronger improvements in 2025, dropping 2.3 points from 2024. Organizations at this size are typically at a maturity threshold where dedicated security teams are being established and enterprise security platforms are being deployed more systematically. The improvement likely reflects that investment beginning to yield measurable results, though the complexity of managing a growing user base and hybrid infrastructure keeps risk elevated relative to smaller organizations.

Organizations with 2,001 to 5,000 employees showed a 3.2-point improvement. Organizations at this scale typically have established security operations centers (SOCs) and dedicated risk management functions. The improvement likely reflects those teams operationalizing more mature vulnerability management and identity governance programs. The remaining risk at 40.7 reflects the inherent complexity of managing large, distributed workforces and multi-cloud environments.

Meanwhile, enterprises with 5,001 to 10,000 employees showed a solid 2.5-point improvement but carry the highest CRI of any size segment in 2025. Organizations of this size operate with network complexity comparable to the largest enterprises but may not yet have the security operations maturity, tooling investment, or staffing depth of organizations above 10,000 employees. This creates a distinct vulnerability profile that security leaders in this segment should pay particular attention to.

The largest enterprises recorded the biggest improvement of any size band, dropping 3.8 points from 2024. This likely reflects the scale benefits of enterprise security programs: dedicated SOCs, automated response playbooks, and significant investment in platforms such as TrendAI Vision One™ that centralize risk management across complex environments. However, at 40.4 the largest organizations still carry elevated risk relative to smaller ones, a reminder that complexity remains the defining challenge at enterprise scale regardless of investment level. The sheer volume of assets, users, and third-party integrations means that even marginal improvements require sustained, coordinated effort across the entire organization, supported by platforms that simplify the complex task at hand.

3

Risk events and detections

Top risk events detected

	Risk event
1	Risky Cloud App Access
2	Stale Microsoft Entra ID Account
3	Virtual Analyzer - Email Risk
4	Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled
5	ZTSA Rule Match - Private Access Control
6	Advanced Spam Protection - Policy Violation
7	On-Premises AD Account with Weak Sign-In Security Policy - Password Expiration Disabled
8	Stale On-Premises AD Account
9	Microsoft Entra ID Account with Weak Sign-In Security Policy - Password Expiration Disabled
10	Virtual Analyzer - Cloud App Risk

Table 1. The top 10 most detected risk events in 2025 (January – December)

Risky Cloud App Access retains the top position, consistent with 2024. The ongoing enterprise cloud migration trend shows no sign of slowing, and organizations continue to struggle with governing which cloud applications employees access.

Stale Microsoft Entra ID Account holds the second spot, suggesting that identity hygiene remains a systemic issue. The persistence of stale accounts in the top two for a second consecutive year signals that remediation efforts are not keeping pace with account sprawl.

ZTSA (Zero Trust Secure Access) Rule Match - Private Access Control, ranked fifth, is a new entrant. Its appearance likely reflects expanded zero trust adoption across the customer base, creating visibility into policy violations that were not previously captured. As organizations move away from traditional perimeter-based security models toward identity- and policy-driven access control, zero trust frameworks are being deployed to govern how users and devices access private applications and resources. The appearance of this event in the top five means that these frameworks are now generating detections at significant scale.

Accounts with MFA disabled and password expiration disabled remain persistent concerns, consistent with 2024's finding that weak authentication policies are widespread.

Virtual Analyzer - Cloud App Risk is new to the list, suggesting cloud application sandboxing detections are increasing alongside broader cloud adoption.

The industry breakdown suggests divergences in primary attack vectors by sector. Mining is uniquely led by email-borne threats (Virtual Analyzer - Email Risk) rather than cloud app access, suggesting targeted spear-phishing campaigns may be a primary vector for that sector.

Healthcare, agriculture, and telecommunications all lead with Risky Cloud App Access, while education is led by the Stale Microsoft Entra ID Account event—pointing to its well-documented identity management challenges across large, distributed user populations.

Top risk events detected in the top 5 industries with the highest average CRI

	Mining	Healthcare	Agriculture	Telecommunications	Education
1	Virtual Analyzer - Email Risk	Risky Cloud App Access	Risky Cloud App Access	Risky Cloud App Access	Stale Microsoft Entra ID Account
2	Risky Cloud App Access	Stale Microsoft Entra ID Account	Data Loss Prevention - Email Violation	Stale Microsoft Entra ID Account	Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled
3	Stale Microsoft Entra ID Account	Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled	Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled	Advanced Spam Protection - Policy Violation	Risky Cloud App Access
4	Malware Scanning - Email Threat	Virtual Analyzer - Email Risk	Advanced Spam Protection - Policy Violation	Data Loss Prevention - Email Violation	Microsoft Entra ID Account with Weak Sign-In Security Policy - Password Expiration Disabled
5	Virtual Analyzer - Cloud App Risk	On-Premises AD Account with Weak Sign-In Security Policy - Password Expiration Disabled	Data Loss Prevention - Cloud App Violation	Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled	Microsoft Entra ID Account with Weak Sign-In Security Policy - Strong Password Disabled

Figure 4. The most detected risk events for each of the top five industries with the highest annual average CRI for 2025 (January – December)

Mining is the only sector in the top five where email-borne threats lead over cloud app access, pointing to targeted phishing campaigns as the primary initial access vector. For a sector with significant operational technology exposure and limited security maturity, a successful email-based compromise carries consequences that extend well beyond the endpoint.

Healthcare's risk event profile reflects its hybrid environment reality: cloud access risks at the top alongside on-premises Active Directory (AD) account weaknesses in the bottom half. This split indicates that identity risk is being managed across two parallel infrastructures. The combination means healthcare organizations cannot address their CRI by focusing on cloud or on-premises controls alone.

Agriculture is the only sector in the top five where Data Loss Prevention violations dominate the top risk events, with both email and cloud app violations appearing alongside Risky Cloud App Access. The pattern signals that sensitive supply chain and operational data is leaving the organization through insufficiently governed channels. MFA-disabled accounts compound the risk, leaving the accounts enabling that data movement inadequately protected.

Telecommunications is the only sector with both inbound email policy violations and outbound data loss prevention failures in its top five simultaneously. This risk profile extends beyond the organization to the customers and critical infrastructure that depend on it. For a sector handling vast volumes of sensitive customer data, the combination of inbound threat exposure and outbound data governance failures is particularly consequential.

Four of the education sector's top five risk events are directly related to account and credential security: stale accounts, MFA disabled, password expiration disabled, and strong password disabled. This makes it the most identity-concentrated risk profile of any sector in the top five. Until account lifecycle management is systematized and automated across its large, constantly rotating user population, reducing the education sector's CRI will remain an uphill challenge regardless of what other controls are put in place.

Top TrendAI Vision One™ misconfigurations

	Misconfiguration event
1	Web Reputation Settings in TrendAI Vision One™ Endpoint Security Not Optimized
2	Device Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized
3	TrendAI Vision One™ Endpoint Security Agent Not Supported
4	Anti-Malware Scanning Settings in TrendAI Vision One™ Endpoint Security Not Optimized
5	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security Not Optimized
6	Predictive Machine Learning Settings in TrendAI Vision One™ Endpoint Security Not Optimized
7	Endpoint Sensor Settings in TrendAI Vision One™ Endpoint Security Not Optimized
8	TrendAI™ Apex One Firewall Settings in TrendAI Vision One™ Endpoint Security Not Optimized
9	Application Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized
10	Behavior Monitoring Settings in TrendAI Vision One™ Endpoint Security Not Optimized

Table 2. The top 10 most detected misconfigurations in 2025 (January – December)

Perhaps the most important insight from the misconfiguration data is not any individual item but the relationship between them. Web reputation, anti-malware, predictive machine learning, endpoint sensors, firewall, application control, and behavior monitoring are not independent controls; they form an integrated, layered defense architecture. When all seven are simultaneously unoptimized across the same customer environments, the resulting security posture is not merely the sum of seven individual gaps. Each weakness compounds the others.

Web Reputation Settings tops the misconfiguration list for the second consecutive year, making it the most persistently neglected configuration across the entire customer base. Given that phishing and malicious web content remain primary initial access vectors, leaving this setting unoptimized creates a significant and unnecessary gap in first-line defense. Its persistence at the top of the list across two years suggests that remediation guidance alone is insufficient; automated enforcement or stronger default configurations may be needed to drive measurable improvement.

An unoptimized Device Control Settings configuration represents a broad and underappreciated data exfiltration and malware introduction risk. In an era where remote and hybrid work has normalized the use of personal devices alongside corporate infrastructure, the absence of properly configured device control policies

creates a persistent blind spot. This misconfiguration appearing at the top of the list for healthcare and agriculture specifically (both sectors with large numbers of connected physical devices) reinforces that device governance is a sector-wide challenge, not just an individual configuration oversight.

Unoptimized Anti-Malware Scanning Settings can mean reduced scan frequency, exclusion of critical file types or directories, or disabled real-time protection—each of which creates meaningful gaps in the ability to detect and block malicious files before they execute.

Smart Feedback enables endpoints to contribute anonymized threat intelligence back to the global threat network of TrendAI™, allowing the broader security ecosystem to benefit from detections made in individual customer environments. When this setting is not optimized, organizations are effectively opting out of the collective intelligence loop—reducing both the quality of threat intelligence available to their own environment and their contribution to the broader community defense network.

Top TrendAI Vision One™ misconfigurations detected in the top 5 industries with the highest average CRI

	Mining	Healthcare	Agriculture	Telecommunications	Education
1	Endpoint Sensor Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Device Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Device Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Suspicious Connection Service Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Non-Optimized Security Agent Password Unlock Settings in TrendAI Vision One™ Endpoint Security
2	TrendAI Vision One™ Endpoint Security Agent Not Supported	TrendAI Vision One™ Endpoint Security Agent Not Supported	Application Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Non-Optimized Security Agent Password Unlock Settings in TrendAI Vision One™ Endpoint Security	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security Not Optimized
3	Suspicious Connection Service Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Predictive Machine Learning Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Non-Optimized Security Agent Password Unlock Settings in TrendAI Vision One™ Endpoint Security	Anti-Malware Scanning Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Behavior Monitoring Settings in TrendAI Vision One™ Endpoint Security Not Optimized
4	Non-Optimized Security Agent Password Unlock Settings in TrendAI Vision One™ Endpoint Security	TrendAI™ Apex One Firewall Settings in TrendAI Vision One™ Endpoint Security Not Optimized	TrendAI Vision One™ Endpoint Security Agent Not Supported	TrendAI™ Apex One Firewall Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Application Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized
5	Application Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Endpoint Sensor Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Application Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized	Device Control Settings in TrendAI Vision One™ Endpoint Security Not Optimized

Figure 5. The most detected misconfigurations for each of the top five industries with the highest annual average CRI for 2025 (January – December)

The mining industry's leading misconfiguration is unoptimized Endpoint Sensor Settings. This is notable given that sensors are foundational to visibility, meaning mining organizations may have significant blind spots in their detection capability.

Healthcare's leading misconfiguration is unoptimized Device Control Settings, consistent with the sector's challenge managing the large number of connected medical devices on its networks.

Agriculture also has Device Control Settings as its top misconfiguration, reflecting the sector's growing IoT and automated machinery footprint noted in the 2024 report.

Telecommunications uniquely leads its top misconfigurations with Suspicious Connection Service Settings, which is appropriate given the sector's network-intensive environment where unusual connection patterns are a key threat indicator.

Education's leading misconfiguration is Non-Optimized Security Agent Password Unlock Settings, a finding consistent with its identity-related risk events and the large, often student-managed device population in that sector.

Top TrendAI Vision One™ cloud misconfigurations

	Cloud misconfiguration event
1	Application Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
2	Firewall Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
3	Log Inspection Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
4	Device Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
5	Anti-Malware Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
6	File Integrity Monitoring (FIM) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
7	Web Reputation Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
8	Intrusion Prevention System (IPS) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized

	Cloud misconfiguration event
9	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
10	Activity Monitoring Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized

Table 3. The top 10 most detected cloud misconfigurations in 2025 (January – December)

Organizations deploying cloud workload protection are not consistently optimizing the settings that make that protection effective, suggesting a gap between deployment and operationalization that persists at scale across the customer base.

Application Control tops the cloud misconfiguration list, governing which applications and processes are permitted to run within cloud workloads and container environments. In cloud environments specifically, unoptimized Application Control Settings are particularly consequential because cloud workloads are often internet-facing, highly automated, and subject to rapid change. This creates frequent opportunities for unauthorized or malicious processes to execute if application governance policies are not precisely defined and actively maintained.

Firewall Settings govern network traffic filtering at the workload level, controlling which inbound and outbound connections are permitted for individual cloud instances and containers. Unoptimized settings in this context can mean overly permissive inbound rules, unrestricted outbound connectivity that enables data exfiltration, or default configurations that prioritize availability over security.

Log Inspection analyzes system and application logs in real time to detect security-relevant events, policy violations, and indicators of compromise within cloud workloads. The significance of this misconfiguration extends beyond the immediate detection gap it creates. Compliance frameworks mandate log collection and analysis as foundational controls, meaning that organizations with unoptimized Log Inspection settings may also be carrying unrecognized compliance exposure alongside their security risk.

Device Control in cloud and workload environments governs the interaction between cloud instances and peripheral or storage devices, including virtual device attachments and storage configurations. Its appearance at rank four in cloud misconfigurations, alongside its rank two position in the endpoint misconfiguration list, suggests that device governance is a pervasive challenge that spans both physical and virtual environments. In cloud workload contexts, unoptimized Device Control Settings can enable unauthorized data transfers between cloud instances and external storage,

create pathways for data exfiltration through misconfigured volume attachments, or allow the introduction of malicious content through improperly governed virtual device connections.

Running unoptimized Anti-Malware Settings in cloud environments means organizations are running cloud workloads without consistent malware scanning across the full workload lifecycle. The challenge is compounded by the speed of cloud operations: in environments where workloads are being spun up through automated pipelines, these settings—when they require manual configuration or are not enforced through policy as code—will consistently fall behind the pace of deployment.

Top cloud misconfigurations detected in the top 5 industries with the highest average CRI

	Mining	Healthcare	Agriculture	Telecommunications	Education
1	Application Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Anti-Malware Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	File Integrity Monitoring (FIM) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Log Inspection Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Predictive Machine Learning Settings in TrendAI Vision One™ Endpoint Security Not Optimized
2	Firewall Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Application Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Log Inspection Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Application Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Web Reputation Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
3	Log Inspection Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Web Reputation Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Firewall Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Device Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Anti-Malware Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
4	File Integrity Monitoring (FIM) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Intrusion Prevention System (IPS) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Smart Feedback Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Intrusion Prevention System (IPS) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized
5	Agent Self-Protection Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Device Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Anti-Malware Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	File Integrity Monitoring (FIM) Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized	Application Control Settings in TrendAI Vision One™ Endpoint Security - Endpoint & Workload Security Not Optimized

Figure 6. The most detected cloud misconfigurations for each of the top five industries with the highest annual average CRI for 2025 (January – December)

The mining industry's cloud misconfigurations are topped by Application Control Settings and Firewall Settings. Gaps in these two controls together mean mining environments may be running unauthorized applications with minimal network-level filtering, a high-risk combination.

The healthcare industry's top cloud misconfigurations lead with Anti-Malware Settings, followed by Application Control Settings, consistent with the sector's challenge of securing a diverse, often legacy device estate.

Agriculture's top cloud misconfigurations lead with File Integrity Monitoring (FIM) Settings; this is significant given the sector's IoT exposure, as FIM is a key control for detecting tampering with operational technology systems.

Telecommunications' top cloud misconfiguration is Log Inspection Settings, a critical gap for a sector where detecting anomalous traffic patterns is paramount.

Top XDR model hits detected

	Model name
1	Possible Disabling of Antivirus Software
2	Hacking Tool Detection - Blocked
3	Registry Run Key Creation Pointing to Files in Temp Location
4	Unknown Threat Detection and Mitigation via Predictive Machine Learning
5	Threat Intelligence Sweeping
6	[Heuristic Attribute] Backdoor File Detection
7	[Heuristic Attribute] Possible OS Credential Dumping
8	File Detections in Windows Directory - Blocked
9	Possible Account Compromise - Atypical Travel
10	Eicar Test File Detection

Table 4. The top 10 most detected XDR model hits in 2025 (January – December)

The 2025 XDR model hits data suggests a threat landscape where attackers are prioritizing defense evasion and persistence establishment above all else.

The rise of Possible Disabling of Antivirus Software to the top of the list marks a significant shift from 2024, where credential dumping topped the overall rankings. This change suggests that adversaries have learned that disabling defenses early in the kill chain is more reliable than attempting to evade them.

Taken together, the top 10 model hits paint a picture of attackers who are systematic, persistent, and increasingly focused on ensuring their foothold is secure before moving deeper into the environment.

4

Vulnerabilities and response data

Top 10 most detected unpatched CVEs

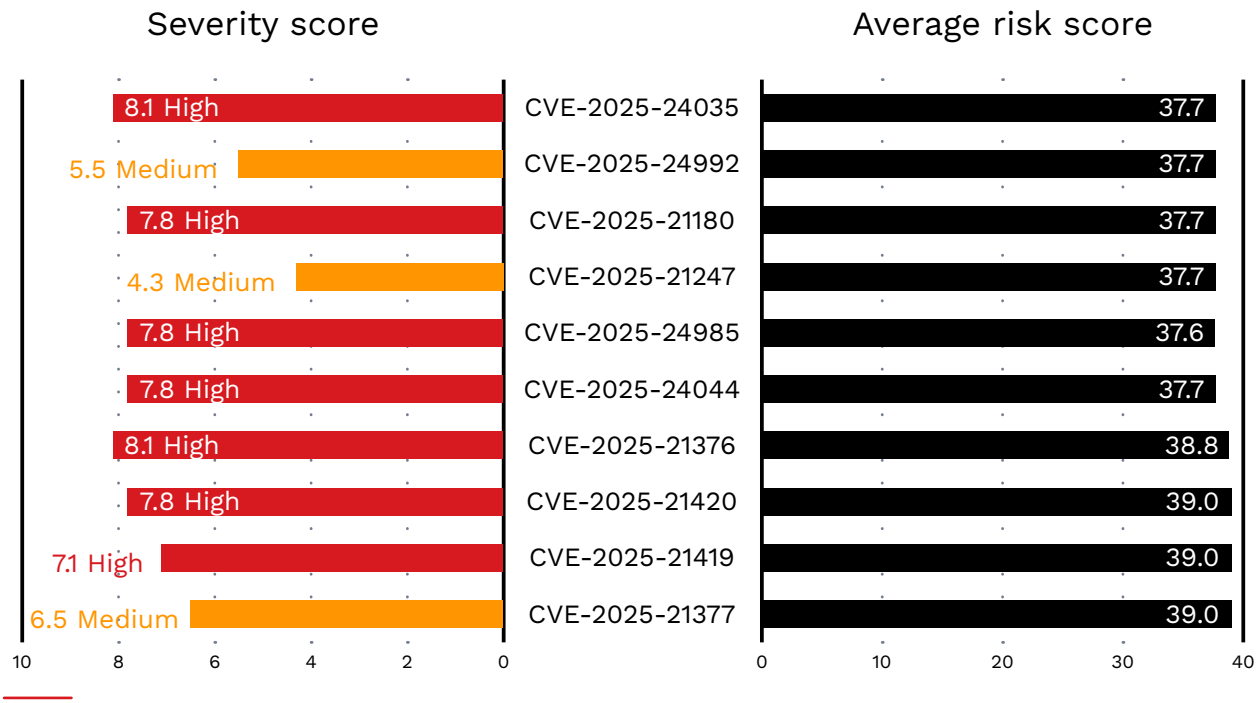

Figure 7. The top 10 most detected vulnerabilities that remain unpatched in customer environments in 2025 (January – December), with their CVSS severity scores from the National Institute of Standards and Technology and their respective average risk scores. Figures have been rounded to one decimal place; ranking reflects unrounded figures.

The list contains a mix of high- and medium-severity scores, but as the analysis below suggests, severity score alone is a poor proxy for actual risk. Three of the 10 CVEs carry medium-severity ratings yet appear in the top 10 most detected unpatched vulnerabilities, a finding that challenges the common enterprise practice of prioritizing patches strictly by Common Vulnerability Scoring System (CVSS) score and deprioritizing anything below a high threshold.

Two vulnerability classes form a natural and dangerous pairing. The first is Remote Code Execution (RCE), covering CVE-2025-24035 (Windows Remote Desktop Services), CVE-2025-24985 (Windows Fast FAT File System), CVE-2025-21376 (Windows Lightweight Directory Access Protocol, or LDAP), and CVE-2025-21180 (Windows exFAT File System), which enable attackers to execute arbitrary code remotely. The second is Elevation of Privilege (EoP), covering CVE-2025-24044 (Windows Kernel-Mode Driver), CVE-2025-21420, and CVE-2025-21419 (Windows system utilities), which provide the escalation pathway from initial access to full system control. The combination of unpatched RCE and EoP vulnerabilities in the same environment is the most dangerous possible pairing, and enterprises must do their due diligence to make sure these vulnerabilities are patched in their systems if they have yet to be.

Three CVEs on the list carry medium-severity scores: CVE-2025-24992 (5.5), CVE-2025-21247 (4.3), and CVE-2025-21377 (6.5), and their presence in the top 10 is one of the most important findings in this year's vulnerability data.

CVE-2025-24992 is a Windows NTFS buffer over-read vulnerability that can lead to local information disclosure, which can enable easier exploitation of the RCE vulnerabilities on this list. CVE-2025-21247 bypasses Windows zone-based content security, allowing untrusted content to appear trusted to the operating system. CVE-2025-21377 discloses NTLM authentication hashes, which can then be used to authenticate as that user on the network.

None of the three medium-severity vulnerabilities is independently catastrophic, but all three become highly dangerous in combination with the high-severity vulnerabilities on this list and the attacker techniques documented in the XDR data. A patch prioritization strategy built on severity score alone will systematically leave these enabling vulnerabilities in place, maintaining the exact conditions that make high-severity exploits more reliable and more damaging. Risk-based vulnerability management requires understanding how vulnerabilities interact, not just how severe they appear in isolation.

Alongside the CVSS severity score, this report includes an average risk score for each CVE, calculated using the TrendAI Vision One™ platform's own risk scoring methodology, which factors in real-world signals such as in-the-wild (ITW) exploitation activity and exploitability, rather than severity alone. These scores are averages across the customer telemetry and are intended to illustrate how risk for a given CVE can diverge from its static severity rating once active exploitation conditions are accounted for. Because actual exposure depends on an organization's specific environment, configuration, and compensating controls, enterprises should treat these figures as directional. They should consult the TrendAI Vision One™ platform to assess the precise risk each CVE poses to their own systems, and patch accordingly.

115 days

Virtual Patch Lead Time

Average, measured from vulnerability discovery (TrendAI™ TippingPoint™, 2025 vulnerabilities)

39 days

Time to Exploit

Average, measured from public disclosure (general vulnerability landscape)

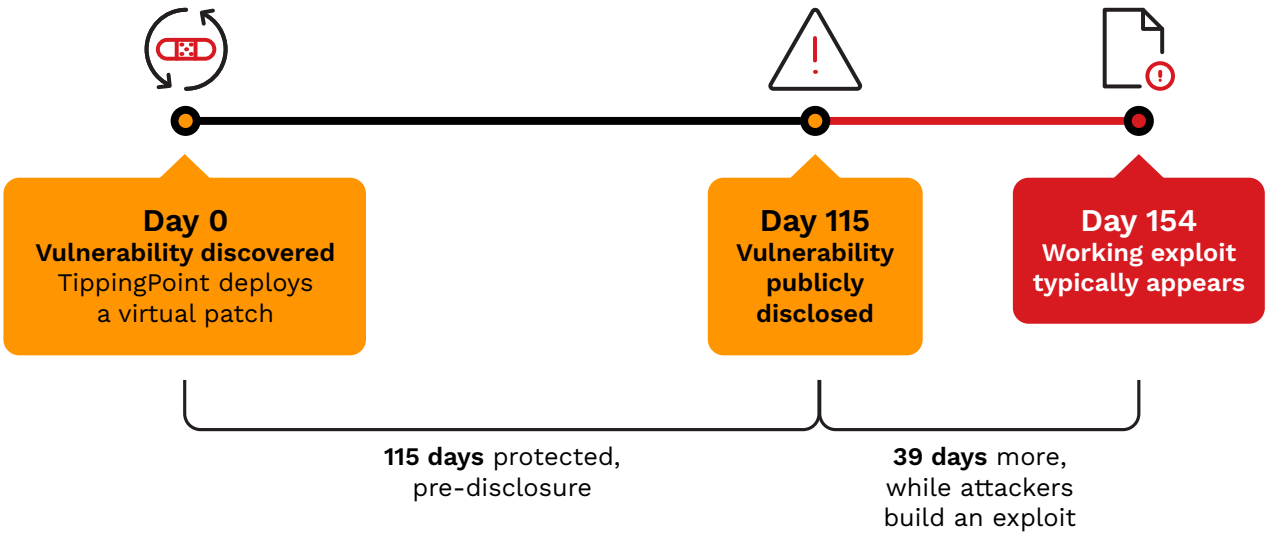

It should be noted that these two numbers don't measure a race to the same finish line but rather measure two different starting points. Virtual Patch Lead Time starts the clock at discovery, when TrendAI™ TippingPoint™ first identifies a vulnerability, often long before anyone else knows it exists. Time to Exploit starts the clock at public disclosure, the moment attackers finally get the information they need to begin building an attack.

Because discovery always comes before disclosure, TippingPoint customers are already covered well before an attacker's clock even starts ticking. On average, that's a 115-day head start with a virtual patch in place, and by the time attackers have a working exploit roughly 39 days after disclosure, TippingPoint customers have been protected for up to 154 days.

That's the core value of virtual patching. It closes the gap between discovery and an official vendor patch, giving organizations a continuous layer of protection while formal remediation works through testing and deployment—no exposure window, no waiting on disclosure to start defending.

The risk, however, is in treating virtual patching as a substitute for formal patching rather than a bridge to it. While the 39-day weaponization window is short, organizations that rely on virtual patching as their primary control rather than a temporary measure could find themselves in an active exploit window if rule deployment runs behind.

This risk is compounded for lower-severity vulnerabilities. Organizations are likely to deprioritize CVE-2025-24992, CVE-2025-21247, and CVE-2025-21377 in both the formal patch queue and virtual patch rule prioritization. This is despite their documented role in enabling the high-severity exploits on this list. The case against patching strictly by CVSS score applies equally to virtual patch prioritization.

5

Attack Path Prediction data

Understanding how an attacker will move through an environment requires more than a list of vulnerabilities or misconfigured controls. It requires modeling the sequences (the chained steps from an initial foothold through lateral movement to a high-value target) that transform individual exposures into realized breaches.

The Attack Path Prediction capability of TrendAI Vision One™ does exactly that. The dataset presented in this section analyzes active attack paths across the customer base to surface the risk events most commonly driving path creation, the asset types most frequently serving as entry points, and the targets that adversaries most often reach at the end of a successful sequence.

Predicted attack path volume scales with organizational complexity, but the events likely to initiate paths and the assets they terminate on can guide how organizations prioritize risk management.

The Attack Path Prediction findings are based on the latest available telemetry at the time of drafting. Because full-year 2025 Attack Path Prediction data was not yet available, this section should be read as a directional view of emerging attack-path patterns rather than a full-year annual measurement.

Top 5 component risk events in Attack Path Creation

Among the most operationally actionable findings in the Attack Path Prediction dataset is the identification of which specific risk events are most frequently present as components at the entry point of active attack paths. These are not simply the most detected events in the environment; they are the events that are actively enabling attack paths to form by providing adversaries with an exploitable starting position.

The distinction matters. An event that appears frequently in detections but rarely appears at the head of an attack path carries a different remediation priority. That priority is highest for events that consistently serve as the initiating condition for a chain of adversary movement.

Rank	Risk event factor	Category	Total attack paths
1	Vulnerability	CVE	2,317,503
2	Potential Brute Force Attack – Password Spraying	Detection	1,135,327
3	Potential Brute Force Attack – Password Guessing	Detection	880,527
4	Behavior Monitoring Detection for Built-in Windows Tools	Detection	539,106
5	Multiple command-and-control (C&C) Connections via Common Protocols	Detection	452,085

Table 5. The top five risk events with the most attack paths mapped by TrendAI Vision One™ (May – June 2025)

Vulnerabilities sit at rank one with 2,317,503 attack paths, about 15% more than the combined total of the two brute-force credential events below it. With the CVE landscape documented earlier in the report, the attack path data now suggests that those unpatched vulnerabilities are not sitting dormant. They are the single largest source of viable attack path entry points across the customer base.

The second- and third-ranked risk events, Password Spraying and Password Guessing, are both brute-force credential attacks, collectively accounting for the initiation of over 2 million attack paths. Both techniques exploit the same structural weakness: accounts that are protected by weak, reused, or nonexpiring passwords, and whose authentication is not enforced through multi-factor controls. This finding is consistent with the identity-related risk events documented in the CRI data, where MFA-disabled accounts and password expiration disabled policies feature prominently across the highest-risk sectors.

Password spraying (attempting a small number of commonly used passwords against a large number of accounts) is particularly effective against certain environments. It succeeds where password policies are inconsistent or where lockout thresholds are set too permissively to avoid disrupting legitimate users. Password guessing targets specific high-value accounts with a larger set of potential credentials. Together, they represent the most reliable initial access vector against identity infrastructure that is not actively hardened. Their position near the top of this list suggests that adversaries are beginning the majority of their predicted attack paths precisely where the CRI data shows defenses to be weakest.

Meanwhile, the fourth- and fifth-ranked risk events suggest what likely follows a successful vulnerability exploitation or credential-based entry. Behavior Monitoring Detection for Built-in Windows Tools captures the use of legitimate Windows utilities

for malicious purposes. By operating within the bounds of trusted system processes, adversaries reduce their visibility to signature-based detection and extend their dwell time within the environment.

Multiple command-and-control (C&C) Connections via Common Protocols at rank five indicates that a significant proportion of active attack paths include an established C&C channel; a meaningful share of environments in the telemetry have passed the initial access stage and reached a point of active adversary presence.

Connecting component events to CRI findings

The component risk events driving attack path creation in 2025 are not independent from the risk events and misconfigurations documented in the CRI dataset; they can be seen as the downstream consequence of those unresolved conditions.

Unpatched vulnerabilities lead the list by a wide margin, consistent with the CVE findings in a previous section: the same Windows RCE and EoP vulnerabilities are the entry conditions generating over 2.3 million attack paths. The CRI data suggests organizations are not closing these vulnerabilities fast enough; the attack path data suggests adversaries can exploit that lag at scale.

The credential attack events at ranks two and three are the consequence of the identity governance failures the CRI data has flagged for two consecutive years. Stale accounts and MFA-disabled credentials create the exact conditions password spraying and password guessing require.

Organizations face two distinct and simultaneously active initial access vectors operating at a significant scale: one technical, exploiting unpatched software vulnerabilities, and one human, exploiting credential hygiene failures. Closing one without addressing the other leaves the majority of the attack surface exposed. The Attack Path Prediction data, read alongside the CVE detection volumes and CRI risk events, points to the same conclusion the report reaches elsewhere—the capability to address these risks already exists in the tools organizations have deployed. The gap is operationalization, and the attack path counts illustrate the cost of that gap in concrete terms.

Entry point and target asset types

Beyond the risk events that initiate attack paths, the telemetry suggests consistent patterns in which asset types are most frequently serving as entry points and which are most commonly reached as the terminal target of a completed path. Understanding

both ends of the attack sequence provides organizations with a more precise basis for control prioritization: hardening high-frequency entry points reduces the rate at which new paths are created, while protecting high-frequency target assets reduces the consequence of paths that do complete.

Top 5 entry point asset types

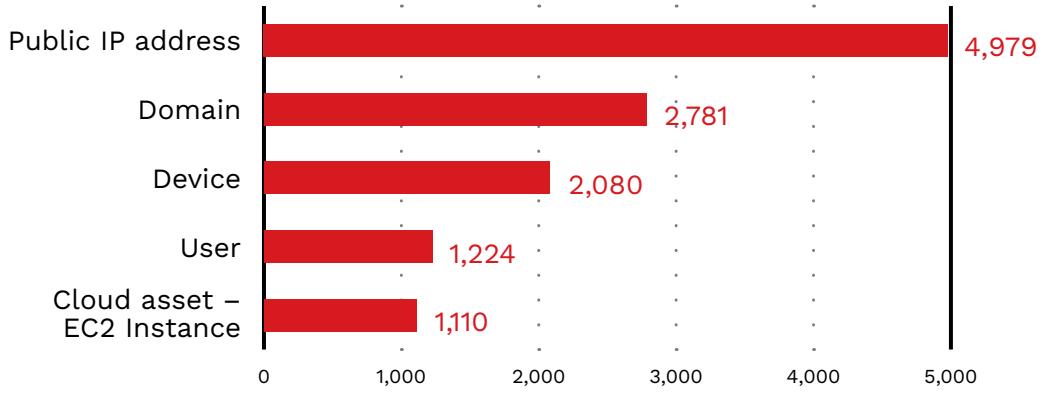

Figure 8. The top five most detected entry point asset types mapped by TrendAI Vision One™ Attack Path Prediction (December 2025 – May 2026)

Public IP addresses are the most common entry point by a significant margin, averaging 4,979 assets per day across the customer base. This finding reflects the perimeter exposure inherent in any internet-facing infrastructure: public IP addresses are the first surface an external adversary encounters and, when insufficiently protected, become the starting point for the attack chain.

Domain assets, at rank two, pair with public IP addresses to define the perimeter exposure dimension of attack path initiation. Where public IP addresses represent the raw network-facing surface, domain assets—Domain Name System (DNS) hostnames attached to those IP addresses—represent the named services sitting on top of that surface: web endpoints, mail servers, virtual private network (VPN) gateways, and administrative interfaces. An adversary approaching an organization's perimeter encounters both simultaneously. The prevalence of domain assets as entry points (2,781 per day) suggests that named services are being targeted as directly as the IP addresses they resolve to. Unoptimized firewall and application control settings documented in the cloud misconfiguration data are particularly relevant here, as these are the controls governing what traffic those exposed services will accept.

Meanwhile, user assets as an entry point suggest that direct account compromise, enabled by the password-based attacks documented in the component risk events section, is a primary mechanism through which adversaries gain their initial foothold.

The appearance of EC2 (Elastic Compute Cloud) Instances at rank five confirms that cloud-native compute assets are now a significant and direct entry point for attack paths, not merely an internal pivot destination. This finding aligns with the cloud misconfiguration data, where unoptimized Application Control and Firewall settings on cloud workloads create direct attack surface in internet-facing cloud infrastructure.

Top 5 target asset types

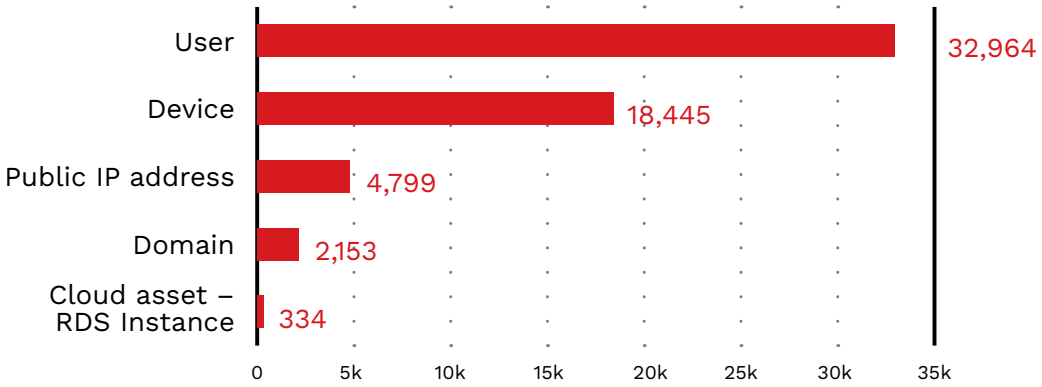

Figure 9. The top five targeted asset types mapped by TrendAI Vision One™ Attack Path Prediction (December 2025 – May 2026)

User accounts are by a commanding margin the most frequent terminal target of completed attack paths, averaging almost 33,000 targeted user assets per day, nearly double the second-ranked asset type. This finding carries direct implications for how organizations should interpret the identity hygiene failures documented throughout this report. Stale accounts, MFA-disabled credentials, and weak password policies are not merely risk events that elevate CRI scores in the abstract. They are the specific conditions that make user accounts the preferred destination for adversary attack chains. An attacker who successfully reaches a high-privilege user account through a completed attack path gains the access needed for lateral movement, data exfiltration, and ransomware deployment.

Device assets at rank two suggest that endpoint compromise remains a primary objective, consistent with the XDR model hit data showing anti-malware disabling and hacking tool deployment as leading adversary techniques.

The presence of RDS (Relational Database Service) Instance at rank five introduces a cloud-native data layer dimension: relational database services are among the most sensitive assets in any cloud environment, and their appearance as attack path targets suggests that adversaries are navigating toward data stores, not just compute nodes.

6

External threat data

Ransomware groups with the greatest number of successful breaches reported in their respective leak sites

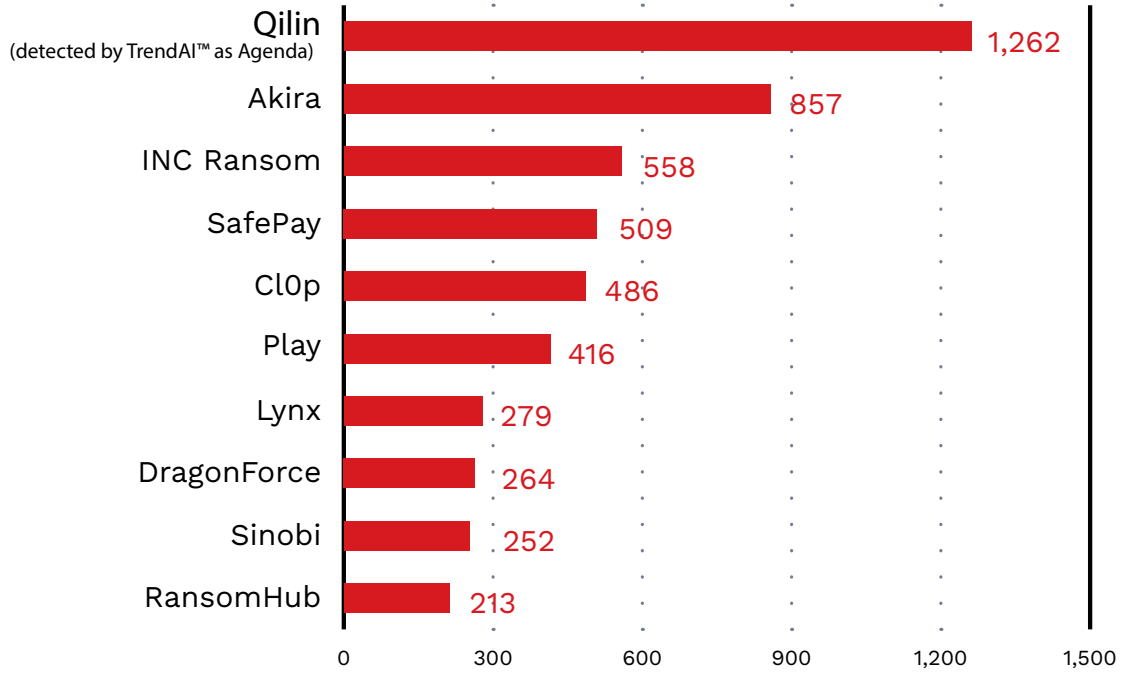

Figure 10. The top 10 ransomware groups with the greatest number of declared successful breaches based on leak site data monitoring by TrendAI™ Research for 2025 (January – December)

The 2025 ransomware landscape based on TrendAI™ Research monitoring of the cybercriminal underground saw a significant reshuffling from last year's findings. "Breach count" here refers to successful ransomware attacks on enterprises that chose not to pay ransom. This means that the actual number of successful attacks might be higher as some companies might have paid the ransom. Total confirmed breach counts across the top 10 have grown approximately 236% compared to 2024, from 1,518 combined victims to 5,096. There are five entirely new groups entering the rankings, two previously dominant groups collapsing, and one group recording a more than 1,200% increase in confirmed breaches year over year. The ransomware ecosystem in 2025 is more fragmented, more active, and more dangerous than the one documented in last year's report.

Qilin ransomware, detected by TrendAI™ as Agenda, climbed from last place in 2024 (with 92 declared breaches) to the most prolific ransomware group in 2025—the single most striking finding in this year's external threat data. A 1,270% increase in confirmed enterprise breaches in a single year suggests that the group has undergone significant operational expansion, likely through aggressive recruitment of affiliates following

the disruption of competing groups. Agenda operates as a ransomware-as-a-service (RaaS) platform, meaning its growth reflects not just the core group's activity but the expanding network of affiliates deploying its tooling across a widening range of targets. Its rapid rise should be treated as a leading indicator of continued escalation in 2026 rather than a peak.

Akira's growth from 106 to 857 confirmed breaches represents a 708% increase and establishes it as the second most prolific ransomware group of 2025. Its rise alongside Qilin suggests that the vacuum left by the disruption of LockBit and the decline of RansomHub has been filled not by a single successor group but by multiple groups expanding simultaneously.

INC Ransom enters the top 10 for the first time with 558 confirmed breaches, establishing itself immediately as a tier-one ransomware group. The group is notable for its targeting of high-value sectors including healthcare and education, and for employing double-extortion tactics that combine data encryption with threatened publication of stolen data.

SafePay is another new entrant, debuting at fourth place with 509 confirmed breaches. Relatively little was publicly known about SafePay heading into 2025, making its sudden appearance near the top of the rankings a significant intelligence gap for the industry.

Lynx also enters the top 10 as a new group with 279 confirmed breaches, adding to the picture of a ransomware ecosystem in which new entrants are achieving meaningful scale rapidly. Lynx has been observed targeting organizations across multiple sectors with a focus on data exfiltration alongside encryption, increasing pressure on victims through the threat of sensitive data publication.

DragonForce also debuts in the top 10 with 264 confirmed breaches, notable for its relatively indiscriminate targeting across industries and geographies. Originally observed as a hacktivist-adjacent group, DragonForce has evolved into a ransomware operator. Its significant breach counts reflect a broader trend of threat actors expanding their operational scope and monetization strategies.

Sinobi enters the rankings with 252 confirmed breaches, representing one of the least publicly documented groups in the 2025 top 10. Its appearance at this scale with limited prior public visibility makes it a priority for threat intelligence teams to track.

Cl0p's appearance at fifth place with 486 confirmed breaches represents a resurgence for one of the ransomware ecosystem's most established actors. Cl0p is historically associated with large-scale exploitation of zero-day vulnerabilities in managed file

transfer and enterprise software platforms, a strategy that generates high victim counts rapidly by compromising widely deployed enterprise tools rather than targeting organizations individually.

For enterprises, the implication is significant. A defensive strategy built around known groups and their established tactics, techniques, and procedures will always lag behind a landscape that is reorganizing this rapidly. The most effective defense is not intelligence about specific groups but reduction of the underlying exposures those groups exploit regardless of who is exploiting them.

External threat geographic and vertical monitoring

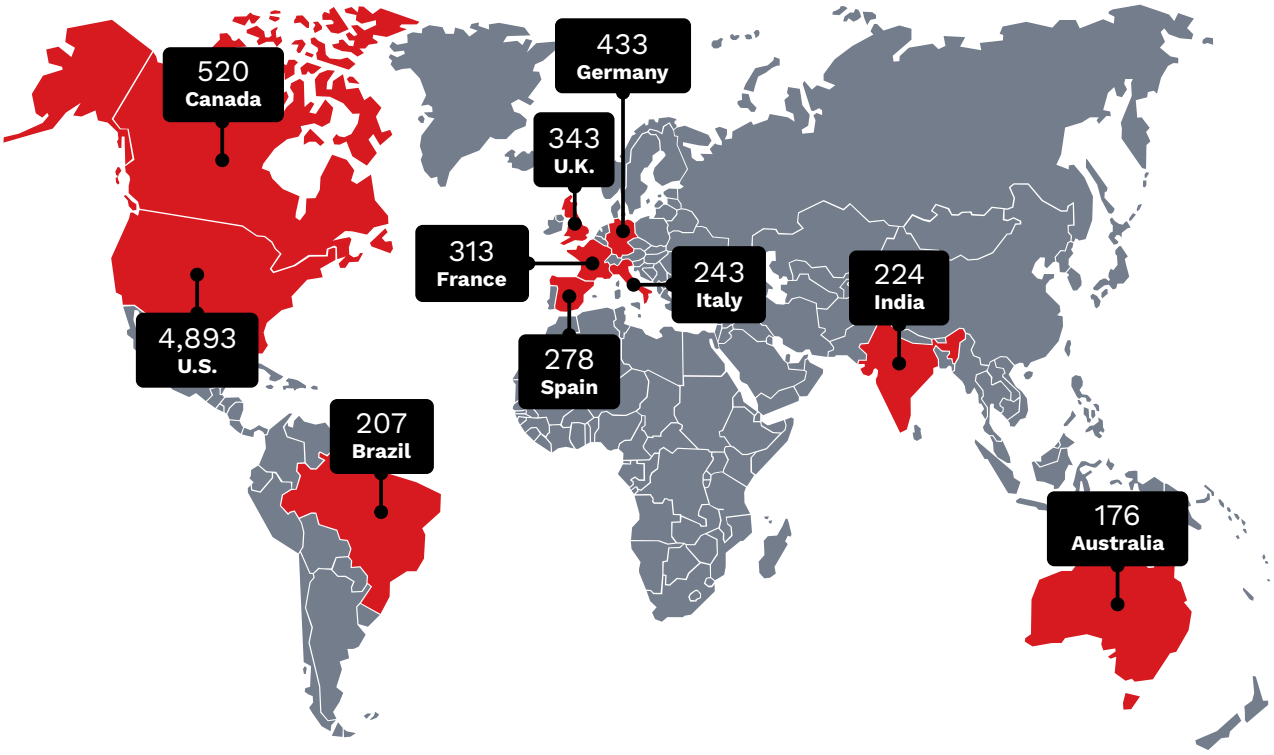

Figure 11. The top 10 countries with the greatest number of declared successful ransomware breaches based on leak site data monitoring by TrendAI™ Research for 2025 (January – December)

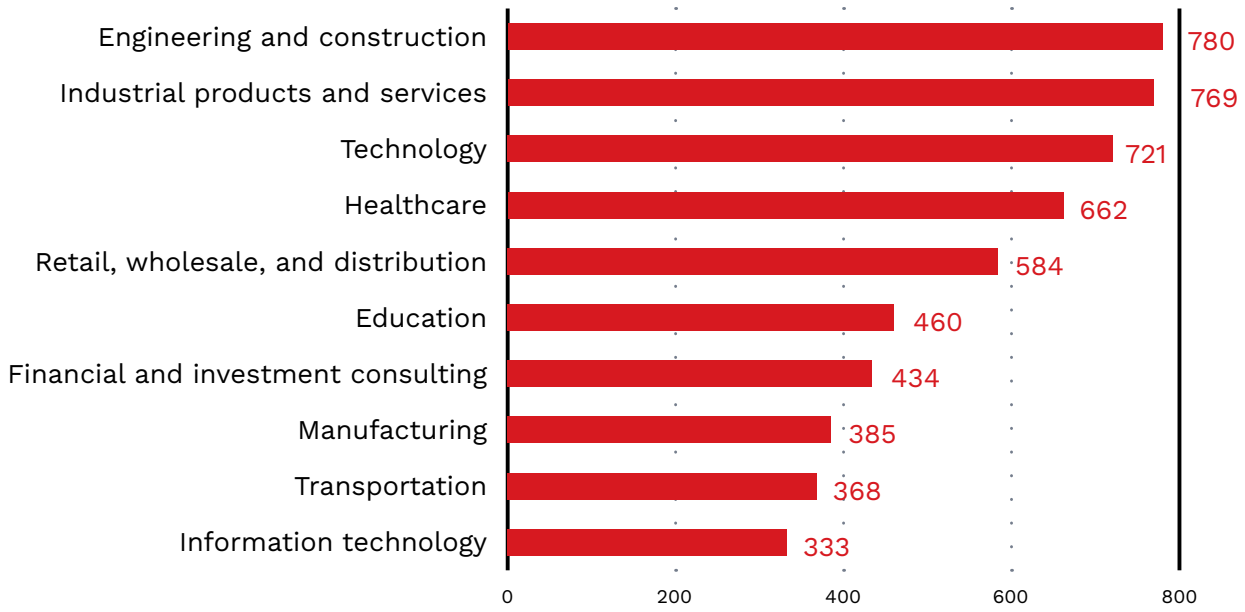

Figure 12. The top 10 industries with the greatest number of declared successful ransomware breaches based on leak site data monitoring by TrendAI™ Research for 2025 (January – December)

7

Conclusion and recommendations

Declining Cyber Risk Index scores suggest that risk-based security management is producing measurable results across the enterprise customer base. Yet the 236% growth in confirmed ransomware breaches, the emergence of five new top-tier threat groups, and the persistence of foundational failures in identity governance, endpoint configuration, and patch management point in a different direction. Together, they suggest that improvement in security posture might not translate into safer environments.

The findings in this report point to a consistent and addressable combination of factors: customer environments continue to run unoptimized security configurations across multiple settings simultaneously, stale and weakly authenticated accounts remain active across the highest-risk sectors, and medium-severity vulnerabilities that could directly enable high-severity exploits are being deprioritized by organizations patching strictly by CVSS score. In each case, the capability to address the risk exists within tools already deployed. The gap is not one of technology but of operationalization.

This year's Attack Path Prediction data makes that gap concrete. Vulnerabilities, password spraying, and password guessing top the list of attack-path initiators, likely feeding on unpatched CVE data and telemetry on unresolved identity hygiene failures. User accounts, the same accounts appearing as top CRI risk events across the highest-risk sectors, are the most frequently reached terminal target of completed attack chains. The exposures organizations are already measuring are the ones adversaries are actively exploiting to build end-to-end attack sequences.

Closing the gap on operationalization requires moving from periodic security activity to continuous risk management. The Cyber Risk Index and other capabilities of TrendAI Vision One™ provide the quantitative foundation for that shift, giving security and business leadership a shared, data-driven basis for understanding exposure, prioritizing remediation, and measuring progress over time. Organizations that have made the most measurable improvement in this year's data did not adopt this model reactively; they operationalized it systematically.

The TrendAI Vision One™ Cyber Risk Exposure Management (CREM) solution provides continuous attack surface visibility across misconfigurations, vulnerabilities, identity risks, and risky events in a unified view. Its XDR capabilities connect that visibility to detection and response across email, endpoint, cloud, and network. Its automated playbooks and AI-guided remediation reduce the operational burden of acting on risk intelligence at enterprise scale. The threat intelligence available within TrendAI Vision One™, used to its full depth, is the difference between a security program that responds to the threat landscape and one that stays ahead of it.

Enterprises must optimize deployed security configurations, enforce identity governance at scale, and adopt vulnerability prioritization that accounts for chained exploit risk. They must also use the full breadth of threat intelligence available in security platforms and solutions that can make risk reduction continuous, measurable, and specific to the environment.

Want more insights like this?

research.trendmicro.com/securitynews

TrendAI™

TrendAI™, the global AI security leader and enterprise business unit of Trend Micro, empowers organizations with full AI visibility and consolidated security that inspires confidence, drives innovation, and eliminates risk.

Trusted by the largest enterprises and governments across 185 countries, TrendAI™ secures the entire organization, from identities to infrastructure to data.

AI Fearlessly.

Learn more: **trendaisecurity.com**