SafeBreach

The SafeBreach 2026 State of the Breach Report

See which threats actually put real enterprises at risk—and where CISOs must focus in 2026 to reduce material impact, not just security noise.

SAFEBREACH.COM

Are we actually protected against the attacks that matter most?

That's the impossible question CISOs face every day. Traditional security metrics—alerts generated, patches applied, or tools deployed—do little to answer this. Defenders need real data about how their organizations (and others like theirs) are actually performing against the real attacker behaviors associated with today's most pressing threats.

This report provides that clarity.

Table of Contents

Introduction

The SafeBreach State of the Breach Report for 2026 analyzes the results of millions of real-world attack simulations executed by large, global enterprises using the **SafeBreach Exposure Validation Platform** over a 12-month period.

As the category leader in Adversarial Exposure Validation (AEV), SafeBreach empowers organizations to continuously and safely validate their security controls against the latest adversarial behaviors, malware families, and CISA-issued alerts.

Our customers include some of the world's largest financial institutions, technology providers, healthcare networks, manufacturers, and critical infrastructure operators, making the SafeBreach dataset one of the richest bodies of empirical security-control-effectiveness data available today.

Throughout 2025, SafeBreach customers executed more than 1.8 million high-fidelity simulations drawn from CISA alerts, nation-state tradecraft, emerging ransomware and infostealers, and industry-specific tactics, techniques, and procedures (TTPs). This report distills that data to answer two critical questions:

What **keeps** CISOs up at night? What **should keep** them up at night?

We will highlight where security programs are performing well, where modern threats continue to evade enterprise defenses, and how trends differ across industries, threat actors, and MITRE ATT&CK techniques.

Our goal is simple: **to provide CISOs and their teams with clear, data-driven insight into their greatest areas of risk and the attacker behaviors most likely to expose their critical assets.**

How We Measure Control Effectiveness

To ground every insight in this report, we start with a simple, consistent framework for measuring how security controls actually perform under real attack conditions.

Every SafeBreach simulation outcome falls into one of three key performance indicators (KPIs), which together form the foundation for all analysis in this report:

PREVENTED

An attacker action was actively blocked by a security control. This represents true control efficacy—where defenses stop malicious behavior outright before impact.

DETECTED

An attacker action was allowed to execute but was logged, alerted on, or otherwise detected by a control. Detection indicates visibility, but not protection; the threat is observed rather than stopped.

MISSED

An attacker action executed without being blocked or detected. Missed activity represents the highest risk outcome, as it reflects attacker behaviors that could progress unnoticed in a real breach.

These three KPIs provide a clear, outcome-driven view of security effectiveness—moving beyond assumptions based on tool deployment, alert volume, or configuration intent. Throughout this report, prevention, detection, and miss rates are used as objective reference points to evaluate controls, architectures, industries, and threat categories.

By anchoring every finding to these KPIs, the State of the Breach Report focuses on what ultimately matters most to CISOs: whether controls stop attacks, merely observe them, or fail entirely when faced with real-world adversary behavior.

Key Insights

The insights below are ordered to reflect what organizations can address most immediately—starting with control coverage and configuration gaps, and progressing toward architectural patterns that shape long-term resilience.

Together, these insights also reflect what organizations choose to validate most frequently—revealing which threats and control layers matter most by industry, and where misalignment between testing focus and real-world risk creates avoidable exposure.

What Industries Test Reflects What They Value—Not Always Where Risk Is Highest

Validation volume varies sharply by industry, signaling perceived priority rather than actual exposure. Most sectors heavily test ransomware and perimeter-focused controls, while under-validating stealth, identity abuse, cloud, and post-compromise behaviors.

A Focus on Both Likelihood & Impact Equals Higher Resilience

Industries that align validation to both ***likelihood*** (what attackers use) and ***blast radius*** (what enables material impact) consistently achieve higher resilience than those optimizing for familiarity alone.

Network Inspection & DLP Controls Lead the Charge

Network Inspection and **Data Loss Prevention (DLP)** controls blocked the most threats, with blockage rates of **~65% and ~70% respectively**; endpoint controls underperformed with a blockage rate of ~53%.

Stealth and Identity Evade

Attacks like ransomware were consistently prevented, but stealthy, identity-driven campaigns (especially **GRU techniques, ~28% missed**) evaded defenses. Gaps cluster around credential abuse and post-compromise movement.

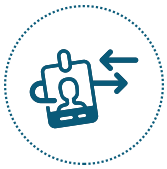

Lateral Movement and Credential Exposure Persist

Segmentation blocked over half of lateral attempts; Retail, SLED, and Healthcare allowed broad traversal. Over **60% of organizations exposed harvestable credentials** (Windows Registry, plain-text), enabling rapid privilege escalation.

AI Exploits Weakness

AI-generated malware was contained, but infostealers showed the lowest blockage. Fragmented/endpoint-heavy architectures performed worst.

Centralized Architecture is Key

Strongest resilience seen in industries with integrated security stacks. Fragmented IT/OT failed regardless of budget/tool count.

Continuous Validation Drives Resilience

Organizations that validated, remediated, and re-validated saw rapid, measurable improvement across all threat categories. Resilience is an operational practice, not a maturity milestone.

Areas of Focus for 2026

To stay ahead of threats in the coming year, the SafeBreach dataset reveals that security leaders need to focus their efforts in four main areas:

1. **Validating controls continuously across multiple environments**
2. **Strengthening identity and data pathways**
3. **Modernizing architectures**
4. **Adopting continuous, empirical testing as part of their security practice**

Mind Your Business: Why Your Industry Matters in Terms of Breach, Risk, & Resilience

Not all industries are created equal when it comes to cyber resilience. The SafeBreach dataset shows a stark truth that CISOs already sense but rarely see quantified: your industry vertical is one of the strongest predictors of how well your defenses actually perform under real attack conditions.

By analyzing the results of millions of simulations run by some of the world's largest enterprises, clear patterns emerge. Some sectors operate with hardened, layered defenses and achieve consistently high prevention rates. Others—often those with complex legacy environments or distributed architectures—struggle to stop even well-known attacker behaviors.

This section breaks down where each industry stands today and why certain verticals outperform or underperform, using real attack-level data rather than surveys, assumptions, or vendor claims.

The Resilience Leaderboard

Industries with modern, centralized security stacks rise to the top when it comes to average blockage rates across all simulations. Those with sprawling OT/ICS systems, distributed endpoints, or inconsistent segmentation fall behind.

Average Blockage Rate by Industry

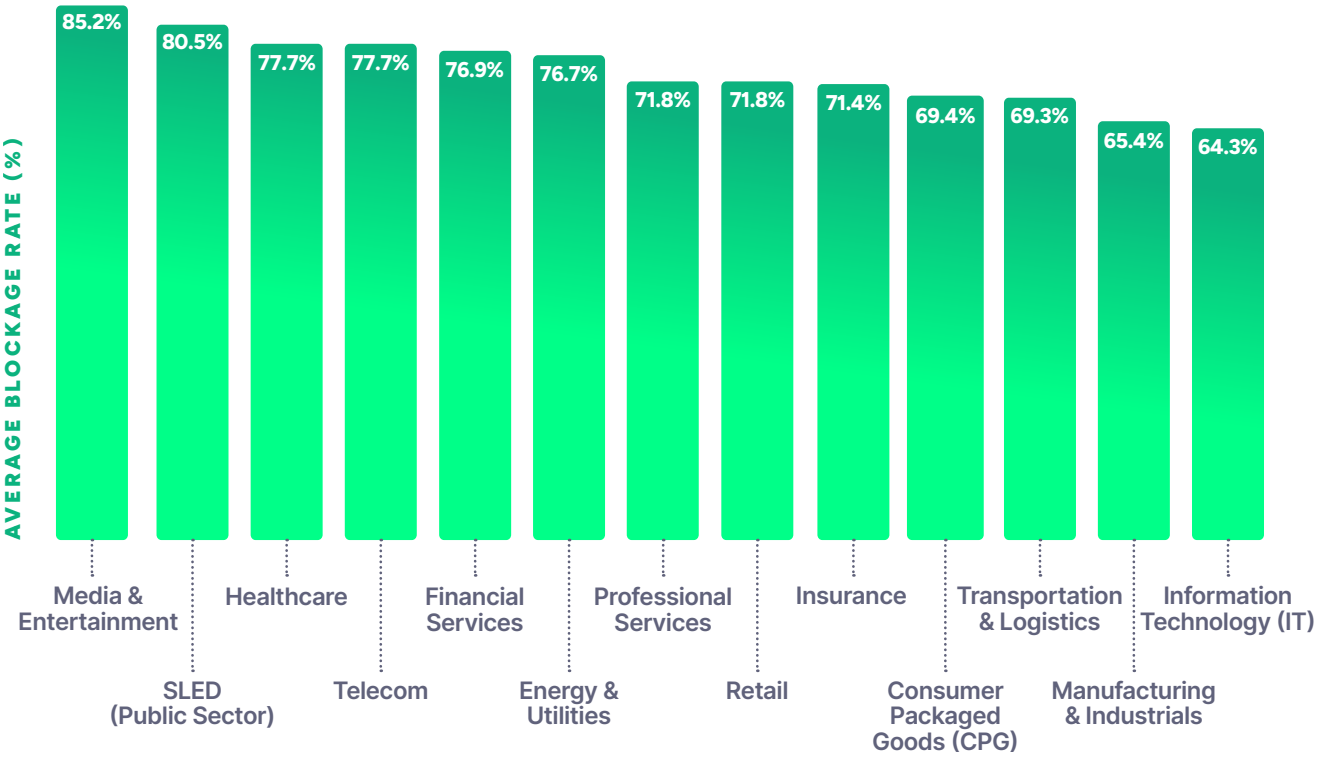

The following leaderboard shows who's keeping attackers out—and who's leaving the door open.

Industry Defense Leaderboard

Top 3 Performers
HIGHEST BLOCKAGE RATES

	INDUSTRY	AVG BLOCKAGE	KEY INSIGHT
1	**Media & Entertainment**	85.2%	Strong layered defenses, likely cloud-heavy
2	**SLED** Public Sector	80.5%	Big spread; some agencies achieving perfect prevention
3	**Healthcare** Tied with Telecom	77.7%	Strong but uneven; likely driven by compliance rigor

Lowest 3 Performers
LOWEST BLOCKAGE RATES

	INDUSTRY	AVG BLOCKAGE	KEY INSIGHT
⚠	**Information Technology (IT)**	64.3%	Surprisingly weak given technical capacity
⚠	**Manufacturing & Industrials**	65.4%	Fragmented control coverage, high attack surface
⚠	**Transportation & Logistics**	69.3%	Distributed Environments, lateral exposure

Control Effectiveness: What Drives Industry-Specific Outcomes

To understand why industries diverge so sharply, we analyzed prevention performance by control category across all simulations.

The first chart shows average blockage rates across sectors, highlighting how differences in architectural maturity and control coverage drive prevention outcomes. The second chart maps each industry's prevention rate against its consistency, revealing how well (or poorly) their control strategies align to the threats they face.

Control Category Performance Analysis

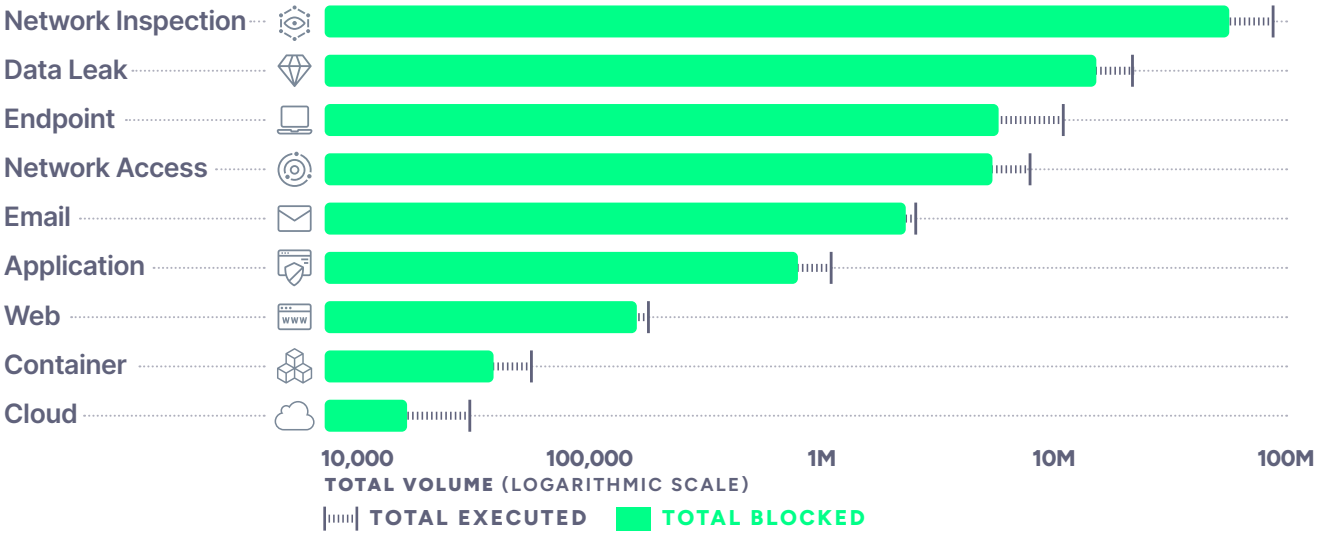

Blockage Efficiency

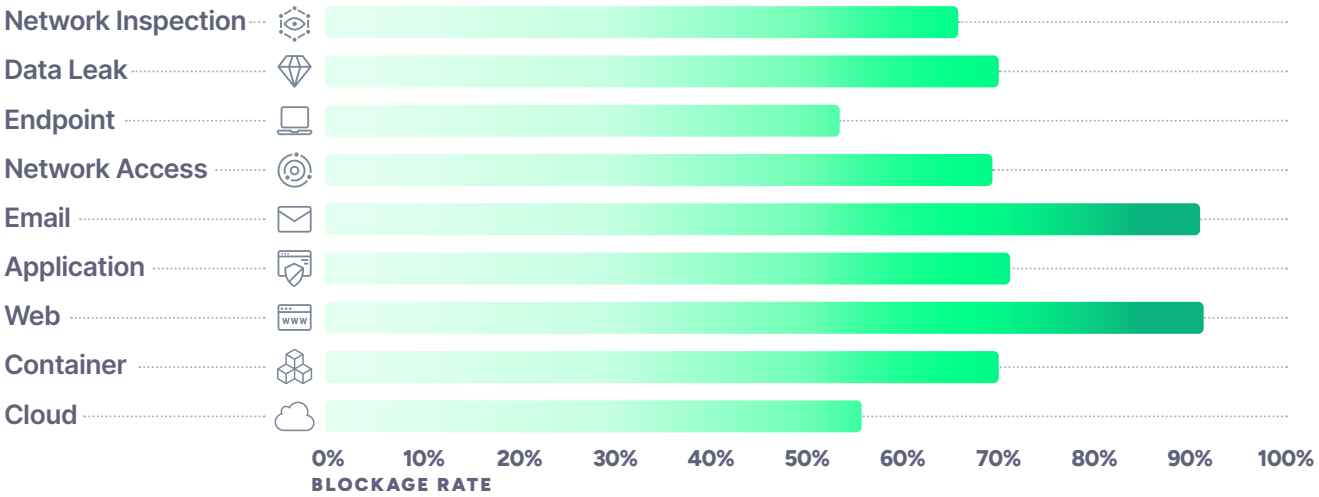

The following quadrant compares each industry's average prevention rate and consistency across attack simulations to illustrate how architectural maturity shapes real-world resilience. Industries cluster, based on how effectively their control strategies align with the threats they face.

Higher detection coverage often reflects delayed prevention rather than stronger security posture. When attacks are blocked early, fewer activities reach later stages where detection alone is required.

Industry x Control Maturity Quadrant

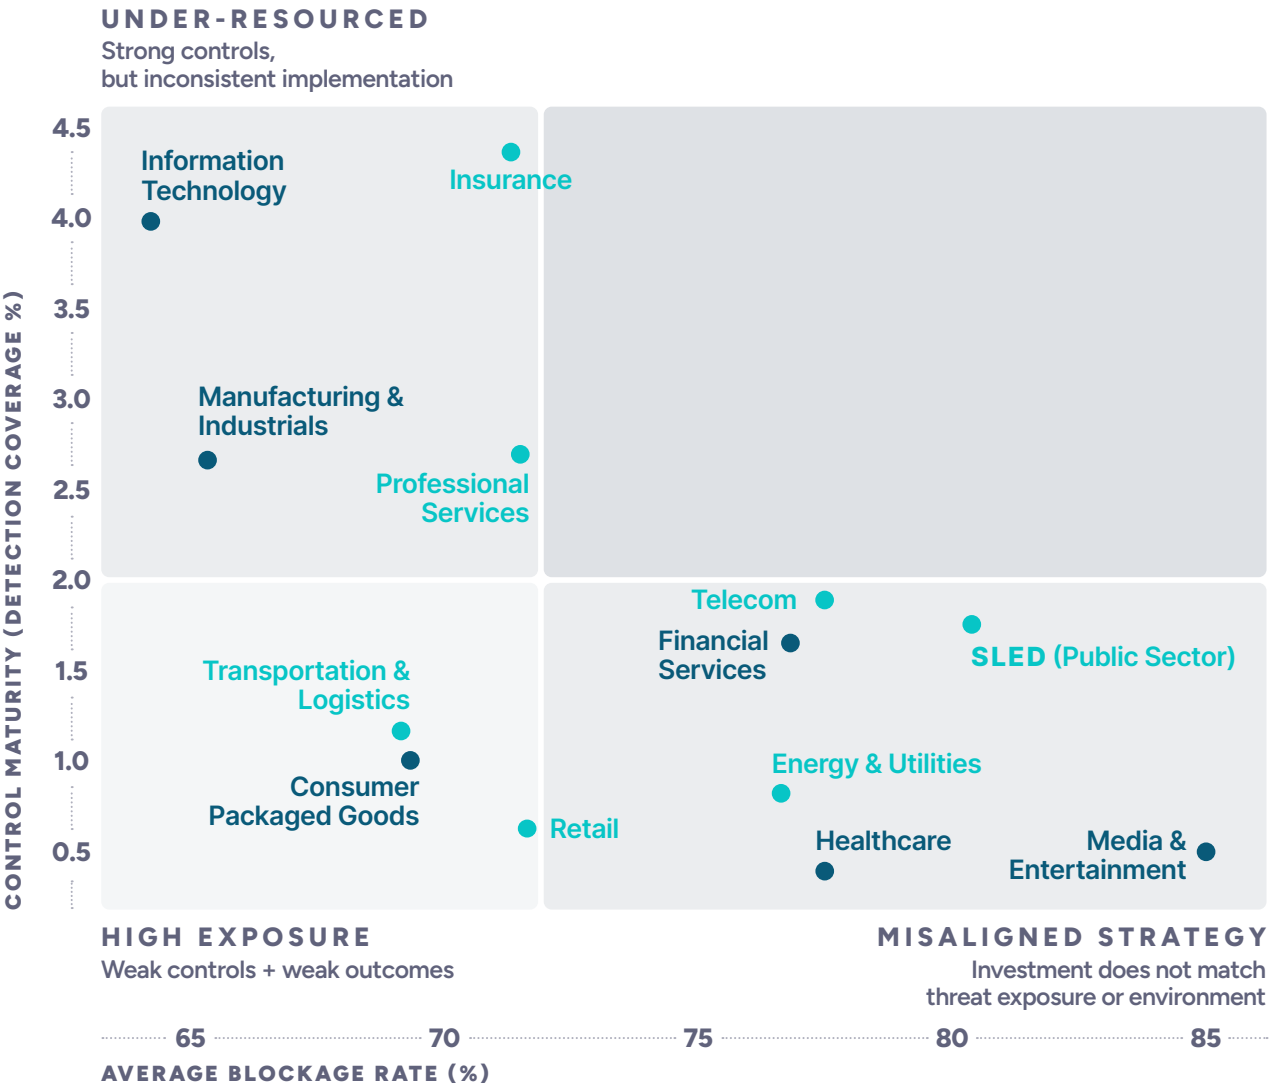

Why Some Industries Win and Others Don't

Industry performance isn't random. It reflects the control layers each sector relies on—and how consistently those controls are validated in real attack conditions. When control-level effectiveness is layered against industry blockage rates, clear structural patterns emerge that explain why some sectors sustain deeper resilience while others face persistent exposure.

THE CONTROLS DOING THE HEAVY LIFTING FOR ENTERPRISE PROTECTION

Network Inspection is the backbone of modern defense. As the number one control by volume—**58.8 million blocked actions**—it delivers outsized impact: **3.5× the effect of DLP** and nearly **9× the effect of endpoint tools**. Industries that invest heavily in network-layer visibility—**Media, Telecom,** and **SLED**—consistently top the resilience leaderboard.

DLP follows as the second-highest blocker at **16.4 million blocked actions**, quietly delivering major value with **~70% blockage**. Data-governance-mature sectors like **Finance, Healthcare,** and **Consumer Packaged Goods** see the strongest benefit, where disciplined data controls translate directly into protection.

Endpoint controls continue to lag. Despite widespread deployment and large budget allocation, endpoints stop **only ~53% of tested attacks**, leaving industries that depend heavily on them—**IT, Manufacturing,** and **Transportation**—firmly in the lower tiers.

The controls an organization validates most often—and how consistently it validates them—offer a clear signal of its operational security.

Control Validation as a Measure of Security Maturity

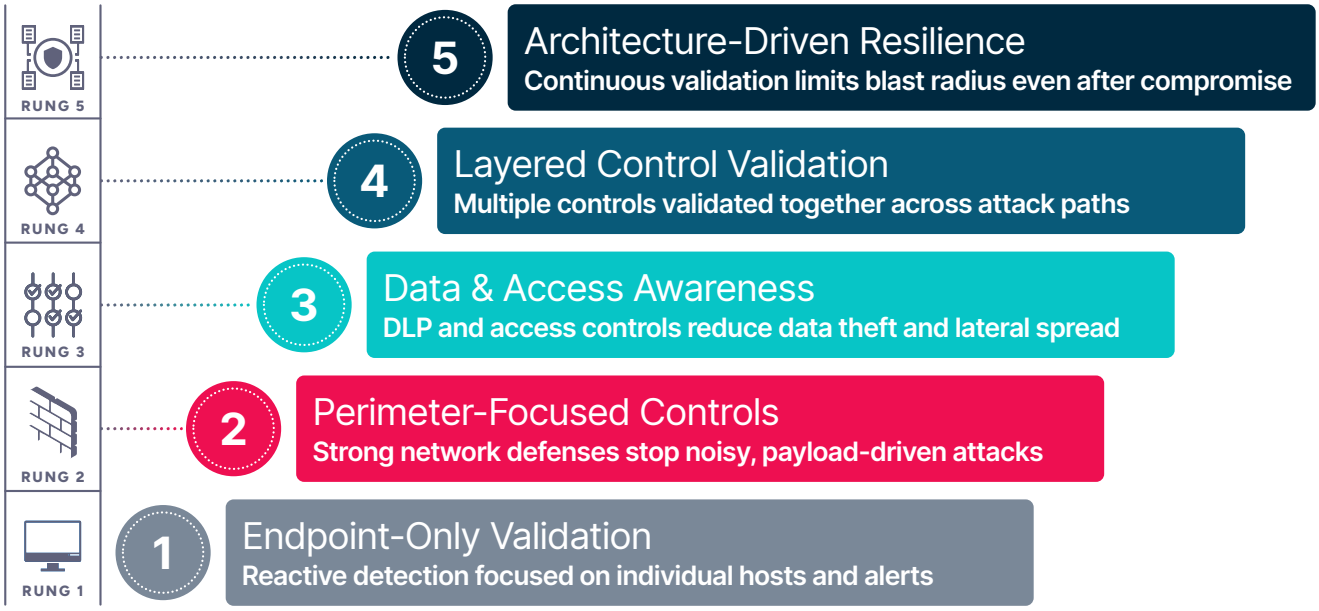

THE FORMULA BEHIND INDUSTRY OUTCOMES

The patterns speak for themselves:

- Strong network inspection + strong DLP = resilient industries
- Endpoint-heavy + fragmented IT/OT = exposed industries

Industries aren't secure or insecure by nature; they're shaped by the control DNA they rely on.

What CISOs Should Ask Themselves

- ☐ Is my industry a top performer or an outlier?
- ☐ Am I over-reliant on low-efficacy controls like endpoint?

CISA Alerts 2025: What Enterprise Validation Tells Us About Real Threat Readiness

Alerts and advisories from CISA flag the most urgent, actively exploited threats hitting enterprises. Within 24 hours of each publication, SafeBreach Labs analyzes the alert, maps any existing attack coverage, and adds new simulations as needed. This rapid turnaround allows organizations to quickly test their real exposure to the threats CISA says matter most—and see whether their controls can actually stop them.

In 2025, SafeBreach customers rigorously validated their defenses against seven major CISA alert scenarios in the SafeBreach platform, ranging from ransomware families to nation-state intrusion campaigns and widely exploited vulnerabilities. Across these scenarios, enterprises executed over 800,000 simulated attacker actions, providing a uniquely empirical view into how well organizations respond when CISA says: “Pay attention.”

The Seven Major CISA Alert Scenarios of 2025

These seven alerts reflect the full spectrum of today's threat landscape: ransomware payloads, multi-stage intrusions, living-off-the-land (LOTL) techniques, authentication manipulation, and supply-chain vulnerabilities. Taken together, they provide a clear signal about what defenders prioritize—and where defenses hold or fall short.

AA25-050A
Ghost Ransomware

Ransomware Actor

53% 37%

Insight: Strong network blockage but applications are bypassed

AA25-141A
Russian GRU Campaign

Nation-State Actor

63% 28%

Insight: Identity and stealth techniques drive enterprise misses

AA25-203A
Interlock Ransomware

Ransomware Actor

60% 40%

Insight: Moderate blocking across the board, but network access is a strong defense

AA25-071A
Medusa Ransomware

Ransomware Actor

85% 15%

Insight: Highly blocked by most controls, but container security is a potential weak point

AA25-141B
Sticky Werewolf Campaign

APT Group

65% 35%

Insight: Network access controls are effective, but endpoint security is a concern

AA25-239A
Chinese State-Aligned Actors

Nation-State Actor

60% 40%

Insight: Application controls are effective, but data leak prevention is a weak area

AR25-261A
Ivanti Endpoint Mobile Vulnerabilities

Vulnerability Exploitation

76% 24%

Insight: Container security is a strength, but data leak and network controls need improvement

Ransomware & Nation-State Campaigns Dominate Enterprise Validation

Simulation volume tells a clear story: enterprises spent 2025 fixated on the threats that hit hardest. Ransomware powerhouses like Medusa and Ghost dominated testing, while nation-state campaigns—from Russian GRU tradecraft to Chinese threat actors and Sticky Werewolf—commanded equally intense scrutiny.

CISA Alerts by Simulation Volume

1. **Medusa**
   583,547 Simulations
2. **Russian GRU**
   70,622 Simulations
3. **Ivanti**
   60,344 Simulations

- **Sticky Werewolf**
  55,204 Simulations
- **Chinese Actors**
  39,178 Simulations
- **Interlock**
  31,503 Simulations
- **Ghost**
  31,216 Simulations

This distribution underscores a consistent trend: CISOs are prioritizing both well-understood payload threats and the stealthy, identity-driven tradecraft associated with sophisticated adversaries.

Enterprises Struggle to Detect Stealth

Across the seven CISA alerts, enterprise defenses performed strongly against high-noise, payload-centric attacks. Medusa, Ghost, and Ivanti scenarios all demonstrated high rates of blocked or prevented activity.

The exception—and the most significant outlier—was the Russian GRU campaign. Here, only ~53% of attacker actions were blocked, and nearly 28% were missed outright. GRU activity consistently bypassed enterprise controls through multi-stage lateral movement, LOTL behavior, credential abuse, and stealthy command-and-control channels.

Top Performing Scenarios

Medusa	85% Blocked
Ivanti	76% Blocked
Sticky Werewolf	65% Blocked
Chinese Actors	60% Blocked
Interlock	60% Blocked
Ghost	63% Blocked

Lowest Performing Scenario

Russian GRU	53% Blocked

This pattern highlights a clear readiness gap: enterprises reliably stop loud attacks but continue to struggle with stealth, identity misuse, and post-initial-access techniques.

How CISA Alerts Impact Industries Differently

The SafeBreach dataset shows that industry architecture materially shapes readiness, with the differences becoming most pronounced in stealth-focused scenarios. Ransomware alerts such as Medusa and Ghost were broadly well-contained across sectors, with high rates of fully blocked activity. But as scenarios shift toward stealth, persistence, or state-aligned tradecraft—particularly GRU and Chinese-aligned campaigns—missed and logged actions rise sharply in industries with fragmented or complex environments, including Manufacturing, Transportation & Logistics, and Healthcare. Industries with more centralized, network-centric architectures maintain stronger containment, even as adversaries pivot to low-noise techniques.

Not all industries appear in every CISA alert scenario, and this is not a data gap—it reflects how organizations prioritize validation based on perceived relevance. Many industries test heavily against ransomware while under-validating stealthier nation-state campaigns. This introduces a critical secondary insight: where validation volume is low, blind spots are more likely to exist, especially as attackers increasingly favor identity abuse, persistence, and lateral movement over loud, payload-centric attacks.

Industries with strong network-centric controls (e.g., Energy & Utilities, Financial Services, and Healthcare) maintained higher levels of prevention across all CISA scenarios. Meanwhile, sectors with distributed or mixed IT/OT environments (e.g., Manufacturing, Transportation & Logistics, and Financial Services) show higher exposure, reflecting gaps in segmentation, identity governance, and endpoint dependency. These distinctions highlight where attackers are most likely to gain traction and where organizations need more targeted validation.

AA25-050A Ghost Ransomware

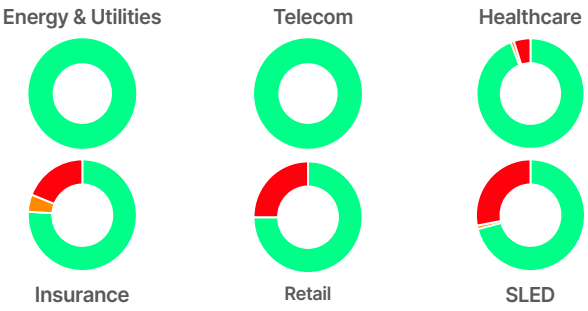

AA25-071A Medusa Ransomware

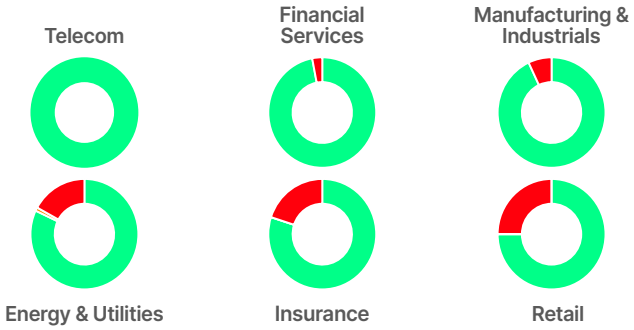

AA25-141A Russian GRU Campaign

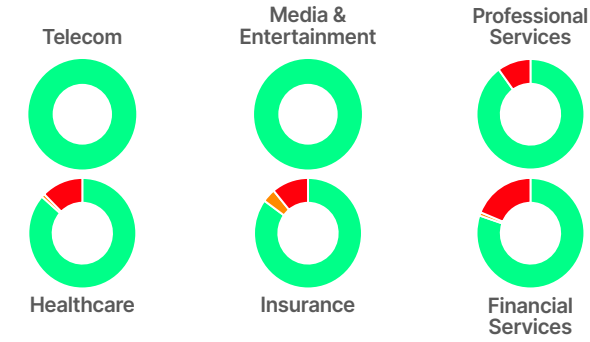

AA25-141B Sticky Werewolf

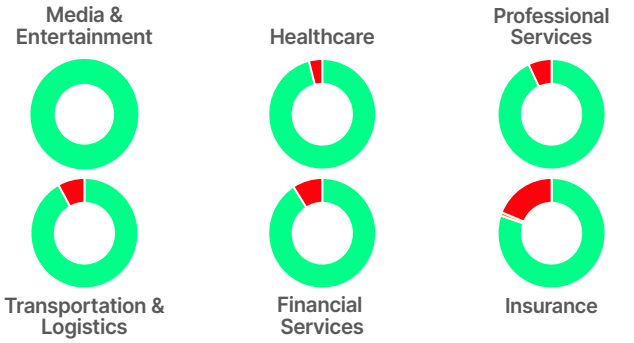

AA25-203A Interlock Ransomware

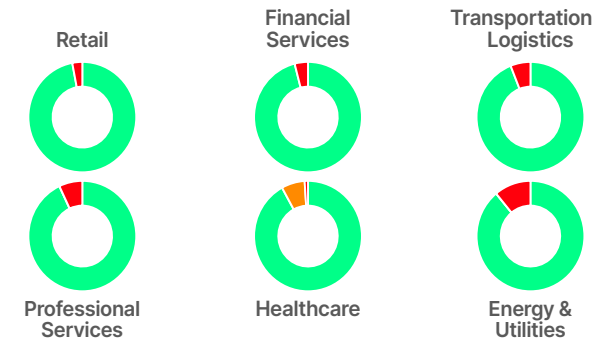

AA25-239A Chinese State-Aligned Actors

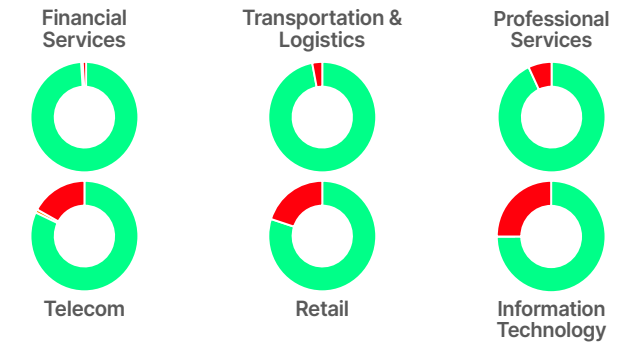

AR25-261A Ivanti Endpoint Mobile Vulnerabilities

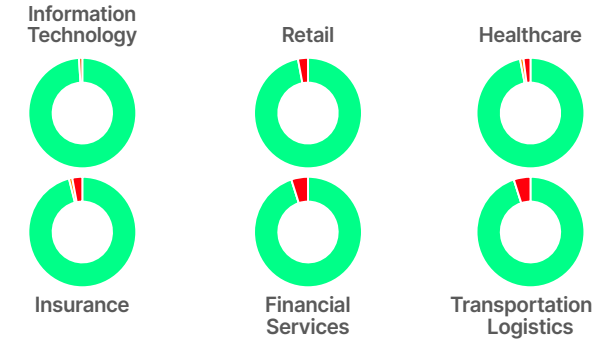

Industry-Specific Exposure Trends

Certain CISA scenarios reveal industry-specific exposure more clearly than others. Here are a few selected industry-specific insights.

HEALTHCARE

Healthcare is good at stopping noisy ransomware but very vulnerable to stealthy, state-aligned campaigns (like those from China) that mimic normal clinical workflows. SafeBreach simulations showed strong prevention against high-noise threats but poor detection of low-noise activity.

This is due to the nature of Healthcare environments: distributed identities, mixed systems, and a security focus on continuity of care over aggressive blocking. Consequently, low-noise persistence, credential misuse, and data-access often resemble routine operations, increasing attacker dwell time and lowering prevention success.

MANUFACTURING

Manufacturing environments underperform in stealth-focused CISA scenarios, especially against state-aligned actors. While ransomware is often contained, multi-stage intrusions (like credential abuse and lateral movement) result in more logged and missed outcomes. This stems from OT-integrated networks' limited segmentation and endpoint controls, allowing attackers to move freely post-access.

TRANSPORTATION & LOGISTICS

Transportation and logistics organizations show high exposure to CISA persistence and identity-driven scenarios. Their distributed infrastructure, reliance on third parties, and mixed IT/OT environments lead to high miss rates when adversaries prioritize stealth. This allows for extended dwell time, even when perimeter defenses block overt ransomware.

What CISOs Should Ask Themselves

- ☐ Are we overconfident in our ransomware readiness—and ignoring the identity-driven, low-noise attacks we actually miss?
- ☐ Where do identity and data pathways remain exposed, and how easily can attackers harvest or abuse credentials in our environment?
- ☐ Does our architecture—especially OT, distributed systems, or weak segmentation—make us inherently more vulnerable than cloud-forward sectors?
- ☐ Are cloud, container, and SaaS attack paths being validated regularly, or are they still untested blind spots?
- ☐ Would we detect or block modern stealth techniques (GRU tradecraft, Storm-0501, AiTM, OAuth hijacking), or would they slip past our controls?

AI-Generated Threats 2025: What Enterprise Validation Reveals About Emerging Attack Techniques

Artificial intelligence is reshaping both the attacker and defender sides of cybersecurity, and attacker adoption is accelerating far faster than many enterprises expected. We now live in a world where generative AI enables even low-skill adversaries to quickly create novel, evasive malicious code.

Toward this end, SafeBreach engineered AI-Generated Scenarios: attacker-style prompts to generate unique infostealer, spyware, and malware variants. The scenarios were then validated across real-world security stacks, where early AI-created samples were missed by half of tested tools.

SafeBreach AI Scenarios

In 2025, SafeBreach customers ran more than 950,000 simulations across AI-generated malware, spyware, and infostealer scenarios, producing one of the industry's first large-scale empirical datasets on control performance against AI-assisted threats.

AI-Generated Malware

Automatically produced variants of traditional malware families — designed for rapid payload delivery, file encryption, host compromise, and execution in seconds without human coding.

AI-Generated Infostealers

AI-crafted tools optimized to extract browser data, session tokens, credentials, personal files, and covertly exfiltrate information — often with minimal user interaction.

AI-Generated Spyware

Surveillance-oriented code generated by AI, capable of recording keystrokes, capturing screenshots, audio input, and monitoring user behavior — offering stealthy, persistent data collection rather than immediate disruption.

The results reveal a clear pattern: enterprises are significantly more prepared for AI-generated payloads than for AI-driven stealth and credential-theft behaviors.

Industry Defense Leaderboard

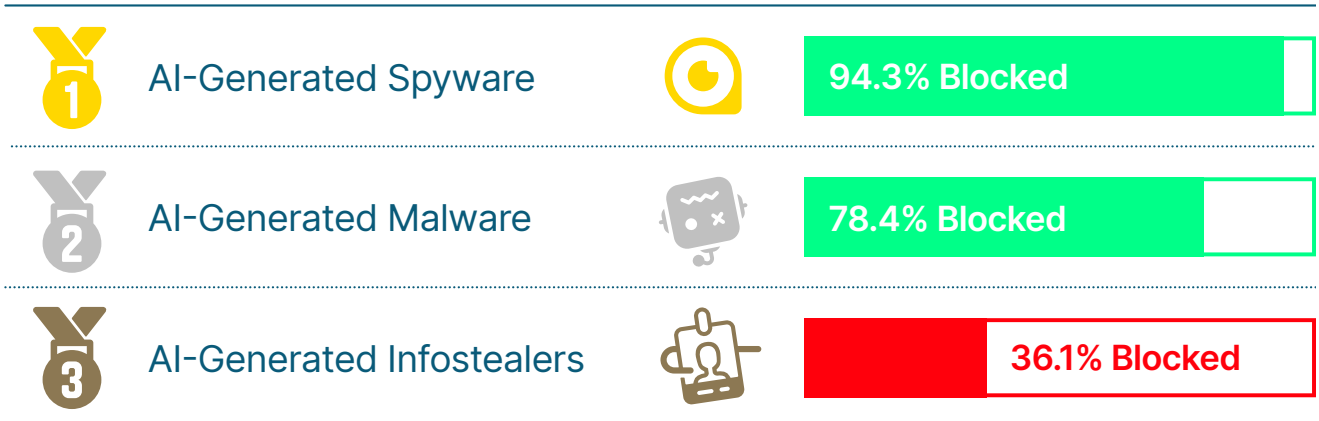

AI spyware shows extremely strong prevention due to mature detection of surveillance-style behaviors. AI malware performs moderately well, suggesting that machine-generated payload variants are still largely within reach of current defenses. AI infostealers, however, remain the most significant exposure point, with the lowest blockage rates and the highest enterprise risk.

Together, these scenarios offer a comprehensive view of how AI is transforming attacker tradecraft across the kill chain—from rapid malware generation to stealthy credential harvesting to long-term surveillance—and highlight the urgent need for security validation that goes beyond signature-based detection and traditional payload blocking.

How Industry Architecture Shapes AI Threat Readiness

This matrix illustrates industry performance against AI-generated attacks, showing average blockage rates and consistency across three AI scenarios. Upper-right industries (e.g., Healthcare, Energy, and Retail) show strong, reliable prevention due to mature network inspection, data governance, and centralized controls. Lower-left industries (e.g., IT Services, Manufacturing, and Automotive) show weaker, inconsistent results, driven by endpoint-heavy strategies and fragmented IT/OT environments.

Blockage Rate x Consistency Quadrant

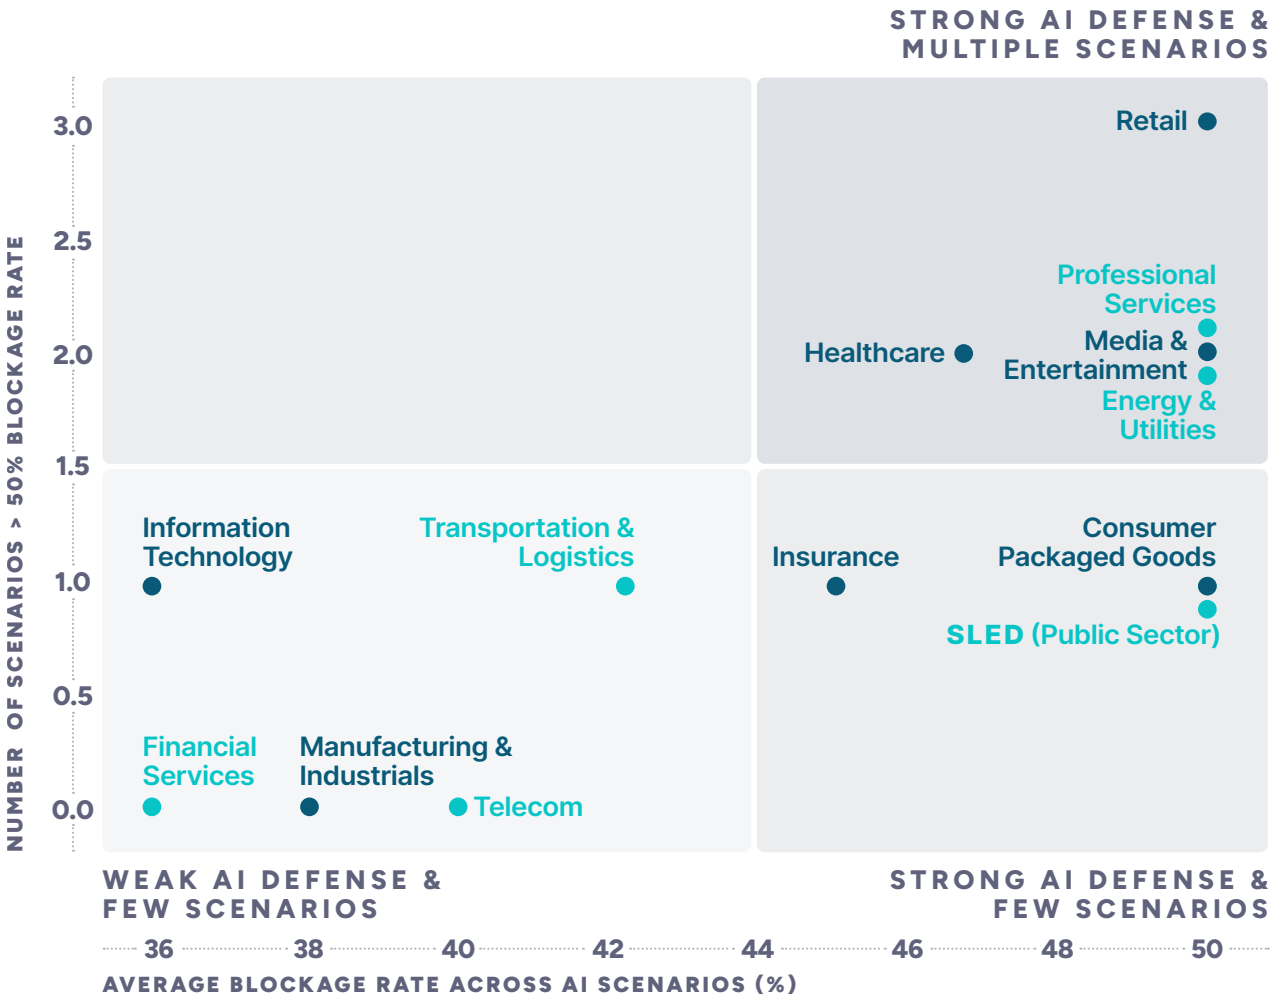

Ultimately, AI-powered attacks intensify existing strengths and weaknesses, widening the gap between sectors with modern, network-centric architectures and those still anchored to legacy footprints.

What CISOs Should Ask Themselves

- ☐ How prepared are we for AI-driven unpredictability?
- ☐ Are we relying too heavily on EDR to stop AI-generated payloads?
- ☐ Where are our identity and data pathways exposed?
- ☐ Does our architecture give us an inherent advantage—or disadvantage?
- ☐ Which environments remain untested—and therefore most vulnerable?

When Breaches Happen: Insights About Lateral Movement with SafeBreach Propagate

SafeBreach Propagate is an attack path validation tool that emulates lateral movement, privilege escalation, and credential harvesting to enable enterprise organizations to understand the potential blast radius of a breach. Data from Propagate simulations shows a consistent pattern across industries: some sectors are substantially more effective at containing attacker movement once an initial foothold is established.

Performance Against Propagation

Industries such as Consumer Packaged Goods and Energy & Utilities block more than half of all propagation attempts, reflecting the benefits of disciplined segmentation, centralized governance, and well-enforced access controls.

Average Blockage Rate by Industry

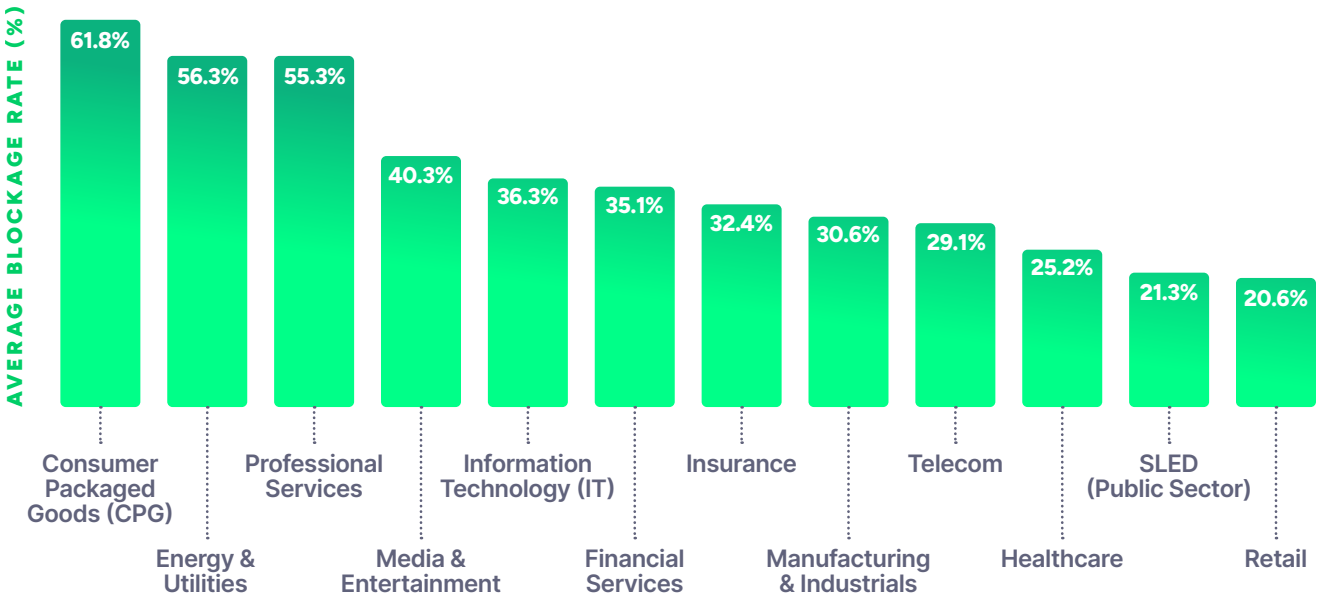

By contrast, sectors with broad, heterogeneous environments—including Retail, SLED, and Healthcare—exhibit significantly lower block rates. These environments often contain hidden pathways, legacy interdependencies, and inconsistent segmentation policies that allow adversaries to move laterally for extended periods before detection.

Here's the bottom line: lateral movement is one of the most accurate indicators of real operational risk. SafeBreach Propagate exposes exactly where attackers are able to traverse the network, turning unseen internal pathways into prioritized remediation tasks. Organizations that identify and close these movement corridors quickly will meaningfully reduce the likelihood that an intrusion becomes a material incident.

Identity Exposure: Credential Harvesting Remains a Critical Weak Point

Across SafeBreach customer environments, identity continues to be one of the most reliably exploitable pathways for attackers. More than **60% of customers** experienced successful credential-harvesting events during testing—showing that even mature enterprises struggle to fully protect identity stores once an adversary gains a foothold.

A deeper analysis of harvested data reveals two dominant—and highly preventable—sources of compromise:

WINDOWS REGISTRY AS A PRIMARY TARGET

Credential harvesting from the Windows Registry emerges as the most common and most fruitful technique. This reflects long-standing architectural realities: cached credentials, service accounts, and application secrets often accumulate in the Registry over time, creating a concentrated, high-value identity store for attackers.

PLAIN-TEXT PASSWORDS PERSIST IN REAL ENVIRONMENTS

While less common, plain-text passwords were found in 1–10% of harvested credentials across affected networks. Even low prevalence represents meaningful operational risk—because each plain-text discovery provides immediate, zero-effort privilege escalation and requires no cracking, guessing, or lateral movement to exploit.

Together, these patterns show why identity remains the decisive battleground for modern security programs. Attackers no longer need sophisticated malware to progress; they simply need access to the right stored credentials. Hardening identity pathways, reducing cached secrets, and continuously validating credential-exposure controls are now core requirements for resilience—not optional hygiene.

What CISOs Should Ask Themselves

- ☐ If an attacker gained a single foothold today, how far could they realistically move before being stopped?
- ☐ Which identity stores would be exposed after initial access—and have we validated that risk?
- ☐ Are we detecting lateral movement early—or simply logging it after the fact?
- ☐ Where do legacy systems or operational constraints create hidden movement paths?
- ☐ Have we prioritized remediation based on attacker reach, not just alert volume?

Conclusion: Building Resilience That Matches Real-World Threats

The data across this report makes clear: real resilience comes from how well security controls perform under authentic attack conditions—not from the number of tools deployed or the volume of alerts generated. Industries with strong network visibility, mature data governance, and disciplined segmentation consistently outperform those relying on fragmented architectures or endpoint-heavy strategies.

CISA alert testing and AI-generated threat scenarios reveal the same pattern: enterprises are increasingly effective at stopping noisy, payload-driven attacks but remain exposed to stealthy, identity-driven techniques—especially AI-powered infostealers and credential-theft campaigns. SafeBreach Propagate data reinforces this reality: once adversaries gain a foothold, lateral movement remains one of the most persistent weaknesses across sectors, allowing small gaps to escalate into significant incidents.

For CISOs, the path forward is clear. Modernize architectures, strengthen identity and data pathways, validate controls continuously across multiple environments, and adopt continuous, empirical testing as a foundational security practice. Organizations that make these shifts will narrow the gap between perceived and actual readiness—and build resilience capable of withstanding the next generation of attacker tradecraft.

SafeBreach enables this transformation by continuously validating controls against real-world attacker behaviors—from CISA alerts to emerging AI-driven techniques—and providing the evidence, prioritization, and remediation insight needed to stay ahead. To benchmark your readiness and uncover your most critical resilience gaps, **talk to one of our offensive security experts today.**

CISO Readiness Checklist

Use this checklist to assess whether your organization is aligned with the resilience patterns surfaced in this report:

VALIDATE CONTROL EFFECTIVENESS

- ☐ Evaluate whether network inspection, DLP, and network access controls are performing as expected
- ☐ Measure endpoint detection performance against stealthy and identity-driven behaviors
- ☐ Validate controls continuously across multiple environments, not just traditional infrastructure

MODERNIZE ARCHITECTURE & REDUCE FRAGMENTATION

- ☐ Prioritize simplification of distributed IT/OT environments
- ☐ Strengthen governance models that centralize control ownership
- ☐ Reduce reliance on endpoint-only security strategies

STRENGTHEN IDENTITY & DATA PATHWAYS

- ☐ Assess exposure to credential theft, session hijacking, and lateral movement
- ☐ Ensure MFA, IAM, and segmentation policies are enforced consistently across all environments
- ☐ Expand visibility into data flows and potential exfiltration paths

ADOPT CONTINUOUS, EVIDENCE-BASED VALIDATION

- ☐ Test regularly against CISA alerts and emerging attacker behaviors
- ☐ Include AI-generated threat scenarios to assess preparedness for rapidly evolving adversary tactics
- ☐ Track both prevention and miss rates to understand true operational resilience

ENGAGE WITH SAFEBREACH

- ☐ Partner with SafeBreach to benchmark your readiness, identify resilience gaps, and continuously validate that your controls can stop real-world attacks to your greatest areas of risk and the attacker behaviors most likely to expose your critical assets.

About SafeBreach

SafeBreach is the leader in enterprise-grade exposure validation, providing the world's largest brands with safe and scalable capabilities to understand, measure and remediate threat exposure and associated cyber risk. The award-winning SafeBreach Exposure Validation Platform combines pioneering breach and attack simulation and innovative attack path validation capabilities to help enterprise security teams measure and address security gaps at the perimeter and beyond. Backed by a world-renowned original threat research team and world-class support, SafeBreach helps enterprises transform their security strategy from reactive to proactive safely and at scale. To learn more about how SafeBreach helps enterprises with end-to-end exposure visibility, visit **SafeBreach.com.**

SafeBreach