



Research by
Cybersecurity
INSIDERS

Zscaler ThreatLabz 2026 **VPN Risk Report**




Table of Contents

Executive Overview	3
Fighting AI with Blindfolds	4
AI-powered Social Engineering: The Credential Theft Multiplier	5
The Response Window Has Collapsed	6
What the Tunnel Hides	7
Lateral Movement Is the Multiplier Once Attackers Get In	8
The Hole In the Firewall	9
One Credential, Full Access	10
Ransomware's Favorite Door	11
The Operational Tax	12
Nation-States Inside the Perimeter	13
The Clock Is the Risk	14
From Reactive to Resilient	15
Zero Trust — The Execution Gap	16
VPN to Zero Trust: Readiness Assessment	17
Predictions For 2026 And Beyond	18
How Zscaler Transforms Secure Access	19
Key Differentiators of Zscaler Private Access (ZPA)	20
Methodology & Demographics	21

Executive Overview

For decades, VPN was the default answer to remote access security — reliable, familiar, and deeply embedded in enterprise architecture. That era is ending. AI has accelerated attack timelines from weeks to minutes, automated credential theft at industrial scale, and given adversaries a speed advantage that human-led defense cannot match. VPN was built for a world where defenders had time to patch, investigate, and respond. That world no longer exists.

Our survey of 822 IT and cybersecurity professionals surfaces a persistent gap: organizations recognize VPN risk clearly, but the perimeter-based access architecture they still depend on cannot contain AI-driven threats that now move in minutes. The remaining question is how fast they replace it.

The VPN architecture itself is the constraint. Faster patching, better monitoring, and tighter policies help at the margins, but none address the underlying exposure VPN creates by design, and AI-driven attacks now exploit that exposure faster than any manual process can close it.

This report examines each risk in detail, quantifies the operational cost, and provides a readiness assessment structured around the CISA Zero Trust Maturity Model to help security leaders measure the gap and prioritize the path from Reactive to Resilient maturity levels. The window to act is measured in the same unit as the threats: minutes.

- **The response window has collapsed.** 79% say the greatest AI-driven risk is attackers weaponizing vulnerabilities faster than patches can be deployed. Only 6% can patch a critical VPN vulnerability within 24 hours; 54% need a week or more. CrowdStrike measured average eCrime breakout time at 29 minutes in 2025, with the fastest at 27 seconds.¹ Attackers operate in minutes. Most defenders still operate in weeks.
- **Defenders are fighting AI with blindfolds on.** 61% report confirmed or suspected AI-enabled attacks in the past twelve months, yet only 5% trust their VPN to detect and stop them. 70% report limited or no visibility into AI-enabled threats traversing VPN connections. With one in five unable to distinguish AI-assisted intrusions from conventional attacks, many organizations will not recognize an AI-powered attack until the damage is done.
- **After compromise, containment fails.** 84% express extreme or significant concern about lateral movement, yet 77% cannot contain it once it starts. Only 11% can restrict a compromised session to a single application, and in one-third of environments a single stolen credential opens the entire network. AI-driven speed makes the initial VPN compromise more likely; once inside, attackers find little to stop them. Ransomware operators exploit this gap systematically — the Akira group collected over \$42 million from VPN-based campaigns confirmed by FBI and CISA.²
- **Encrypted tunnels shield the attacker.** Sixty percent inspect a quarter or less of encrypted VPN traffic. Fifty-two percent describe their VPN as a transport layer with limited or no inspection capability. Malware delivery, command-and-control, and data exfiltration all ride the same encrypted tunnel as legitimate traffic, indistinguishable and uninspected. Attackers don't need to build covert channels when the VPN already provides one.
- **Operational friction is eroding security.** Seventy-three percent say VPN demands more effort than modern alternatives. Forty-five percent report significant or major productivity impact. The predictable result: 63% say users intentionally bypass VPN controls to reach applications faster. When the sanctioned path is the slowest path, the VPN generates the exposure it was designed to prevent.
- **Zero trust adoption is accelerating, but the gap remains.** Eighty-four percent are planning or transitioning to zero trust, up from 78% two years ago. Fifty-one percent report declining VPN usage. The drivers map directly to this report's findings: lateral movement elimination (67%) and reduced operational overhead (62%). Hybrid VPN environments will persist for years, extending exposure during the coexistence period.

1. CrowdStrike, 2026 Global Threat Report, February 2026. 2. FBI and CISA, Joint Cybersecurity Advisory: Akira Ransomware, April 2024; updated reporting through 2025.



Fighting AI with Blindfolds

In this report, “AI-enabled attack” refers to attack activity where respondents observed AI being used to increase speed, scale, or sophistication — for example automated reconnaissance, AI-generated social engineering, rapid exploit adaptation, or evasion techniques. The classification reflects what defenders observed and attributed in their environments, not laboratory-grade certainty.

AI-assisted attacks have been building for years, but the last twelve months pushed them into day-to-day reality. 18% confirmed AI-enabled attacks in the past twelve months and another 43% suspected them. Together, that means 61% report confirmed or suspected AI-enabled attacks. Only 5% trust their VPN to detect and stop AI-enabled threats.

One in five organizations cannot distinguish an AI-assisted intrusion from a conventional attack. That 19% means the incident figures reported here are almost certainly an undercount.

Defensive readiness has not kept pace. Over half say legacy VPN infrastructure blocks integration of AI-driven security tools entirely, which is why only one in four has managed to deploy AI-powered monitoring. 79% fear AI will weaponize vulnerabilities faster than they can patch, yet only 5% trust their VPN to detect and stop AI-enabled threats when patching fails. That 74-point gap between fear and readiness defines the defensive crisis this report measures.

Closing that gap requires moving inspection and detection inline at the access layer, where defensive AI can match the speed of modern automated attack.

Every quarter that gap stays open, the advantage tilts further toward the attacker.

In the past 12 months, has your organization experienced or detected any attacks you believe were AI-enabled or AI-assisted?




Figure 1: AI Attacks Are Already Here

74 pts The gap between fear and readiness
79% fear AI exploit speed vs. 5% trust VPN to stop it

Do you believe your current VPN infrastructure is capable of inspecting and stopping AI-generated attacks?




Figure 2: VPNs Can't See What's Coming



AI-powered Social Engineering: The Credential Theft Multiplier

Which AI-enabled attack techniques are organizations most concerned about targeting their VPN infrastructure? The clear leader is AI-generated phishing or social engineering to steal VPN credentials (63%)—ahead of automated vulnerability scanning and exploitation at scale (57%) and adaptive attacks that evade traditional detection signatures (49%). That ordering matters: it signals that for many adversaries, the fastest path to VPN compromise isn't "breaking in," but logging in.

AI raises both the speed and success rate of social engineering. Pretexts can be customized in seconds, lures can be iterated continuously, and targeting can expand from executives to any user with remote access. In VPN environments, the outcome is the same: once an attacker obtains a valid credential, the VPN often treats that identity as trusted and extends broad connectivity into internal resources—turning a single successful phish into a high-confidence entry point.

Recent credential-theft campaigns show how rapidly VPN phishing is evolving. Since mid-January 2026, researchers reported a threat actor using SEO poisoning to lure VPN users to trojanized installers that display a fake login prompt to steal enterprise VPN credentials, then route victims to the legitimate site to reduce suspicion.³

This is why credential theft remains one of the most structurally dangerous VPN risks. The perimeter-style model authenticates once and then grants access based on that moment in time. Anyone holding a valid credential can inherit the user's reach—often beyond what the role requires, and often without continuous verification of device posture or behavioral signals. In that context, AI-powered social engineering isn't just another technique; it's the multiplier that makes credential compromise reliable, repeatable, and scalable.

Which AI-enabled attack techniques are organizations most concerned about targeting their VPN infrastructure?




Figure 3: Social engineering is the most-feared AI-driven attack technique targeting VPNs

One phished credential. Full access. To break that chain, authentication must shift from a one-time perimeter event to a persistent control—continuously evaluating identity, device posture, and behavior on every request, not just at login.

3. Microsoft, 2026. Storm-2561 uses SEO poisoning to distribute fake VPN clients for credential theft. March 2026.



The Response Window Has Collapsed

Every VPN vulnerability starts the same clock. But can defenders patch before attackers exploit? In 2026, that race is over for most organizations.

In January 2025, a critical Ivanti VPN zero-day (CVE-2025-0282) entered active exploitation before any patch existed; CISA and Mandiant confirmed that post-exploitation malware survived factory resets and firmware updates.⁴

That was one incident. By year's end, three Cisco ASA/FTD zero-days had been exploited by a state-sponsored actor over five months, a WatchGuard Firebox zero-day (CVE-2025-14733) had exposed 125,000 devices globally, and Arctic Wolf had confirmed a sustained SonicWall SSL VPN campaign with short intervals between initial access and ransomware encryption. The cadence is no longer episodic; it is continuous.

Only 6% of organizations can deploy a critical VPN patch within 24 hours. The majority, 54%, need a week or more. The CrowdStrike 2026 Global Threat Report measured average eCrime breakout time at 29 minutes in 2025, with the fastest observed at 27 seconds.¹ For many organizations, compromise occurs before their patch process clears the first approval gate.

Practitioners call it the CVE of the week: wake up to a CVSS 9.8 alert, scramble to deploy, do it again six days later. 56% rank patching as their top operational challenge, and 61% acknowledge that vulnerabilities are weaponized faster than fixes deploy. The defenders who spend the most time patching are also the ones who know it will never be fast enough.

Zscaler ThreatLabz analyzed 411 VPN CVEs over five years and found 82.5% growth in annual volume, with 60% of the most recent year rated high or critical. The vulnerability pipeline is structural, not episodic.

When the initial foothold is the access layer itself, compensating controls sit downstream of the compromise. A patched concentrator prevents the foothold, but once exploitation succeeds, containment depends on segmentation and inspection capabilities that this survey shows are weak across most environments. As long as remote access depends on a customer-managed appliance, every organization remains one unpatched CVE away from compromise. Removing the VPN concentrator from the equation eliminates one of the most exposed patch surfaces in remote access.

From the moment a critical VPN vulnerability is publicly disclosed, how long does it typically take your organization to fully deploy patches across all VPN appliances?




Figure 4: Patching Takes Weeks. Breakout Takes Minutes

What challenges does your organization face with patching VPN vulnerabilities today?




Figure 5: Why Patching Never Catches Up

Multiple responses allowed

4. CISA, Emergency Directive 25-01: Mitigate Ivanti Connect Secure Vulnerabilities, January 2025; Mandiant, CVE-2025-0282 Analysis, January 2025.
1. CrowdStrike, 2026 Global Threat Report, February 2026.



What the Tunnel Hides

What percentage of the SSL/TLS (encrypted) traffic flowing through your VPN is your organization currently able to fully inspect for malware and threats without degrading performance?




Figure 6: The Encrypted Blind Spot

A credential-stuffing bot, a C2 beacon, and a 3:00 AM exfiltration disguised as a file sync all cross the tunnel looking identical to legitimate traffic. The technology built to protect data in transit now provides cover for every stage of an attack.

Inspecting traffic inline at the point of access restores the visibility and control that encrypted tunnels strip away.



52% say their VPN is just a pipe with no inspection capability

The same encrypted tunnel that protects a legitimate session also shields every attack that crosses it. Whether the appliance is fully patched or not.

VPN was engineered to create encrypted tunnels, and at that task it performs exactly as designed. Encrypted transport has since become the primary delivery mechanism for modern threats, and VPN is rarely configured, or in many cases capable, of inspecting what the tunnel carries. Attackers embed payloads in SSL/TLS sessions, run command-and-control frameworks like Cobalt Strike through encrypted connections, and exfiltrate data through the same tunnels that carry legitimate traffic. The result is a blind spot that grows with every encrypted session. Adversaries can hide inside “legitimate” encrypted traffic and adapt their techniques faster than legacy controls can detect what’s moving through the tunnel.

In January 2026, Huntress documented an attack chain where a compromised SonicWall VPN provided initial access, enabling the attacker to pivot through to VMware ESXi hypervisors using an exploit toolkit developed over a year before disclosure.⁵

5. Huntress, SonicWall VPN to ESXi Attack Chain Analysis, January 2026.

The entire chain crossed encrypted tunnels that no inspection layer examined. When 52% describe their VPN as a transport layer with limited or no inspection capability, they are describing what the technology was built to do: systems performing according to design, in an environment that design never anticipated.

One-third of organizations inspect zero encrypted VPN traffic for threats. Add the 27% examining less than a quarter, and 60% of enterprises run remote access with minimal or no visibility into what crosses it. Only 8% can inspect virtually everything. 83% rank ransomware through VPN tunnels as a top concern, yet 60% barely look at what’s inside them. The gap is even wider for AI-enabled threats: 70% report limited or no visibility into AI attacks traversing VPN connections. The organizations with the most to fear have the least ability to see it coming.



Lateral Movement Is the Multiplier Once Attackers Get In

Initial access is rarely the end of the story, it's the beginning of expansion. With VPN-based remote access, a single authenticated session often lands a user on the internal network, turning one compromised credential into broad reach across systems and services.

The report shows how limited containment still is in most environments: only 11% of organizations can restrict a compromised session to a single application. In 32% of environments, a single stolen credential opens the entire network and 14% can't describe their own blast radius.

Nearly half of enterprises can't confidently predict what an attacker could reach after login.

Segmentation gaps make lateral movement easier to operationalize. 56% have no per-application segmentation through VPN, and 24% provide entirely open access after authentication. The result is a structural mismatch: defenders may detect an intrusion, but lack the access controls to contain it quickly.

That tension shows up clearly in priorities: 81% flag lateral movement as a top concern, yet 77% lack confidence in containing it. When remote access grants network-level connectivity instead of app-level containment, lateral movement becomes the multiplier that turns one compromise into many.

Organizations know lateral movement is a major risk, but most still lack the controls to contain it after login



81%
say lateral movement is a top concern



77%
lack confidence containing it once it starts



The Hole In the Firewall

The perimeter was supposed to be the hard boundary. For most organizations, VPN turned it into an open door, and the uninspected traffic documented earlier flows straight through it.

Networking teams have grown vocal about the irony: you spend millions on firewalls, then punch a giant, persistent hole in it for VPN to listen on. Every VPN session passes through the perimeter, places the remote user directly on the internal network, and bypasses the inspection layers the organization spent years building.

Only 11% can restrict a compromised session to a single application. In 32% of environments, a single stolen credential can open the entire network. Another 14% cannot describe their own blast radius. Nearly half of all enterprises cannot predict what an attacker would reach after login.

56% have no per-application segmentation through VPN, and 24% provide entirely open access after authentication. Almost everyone who recognizes the threat also lacks the means to address it: 81% flagged lateral movement as a top concern, yet 77% lack confidence in containing it.

Site-to-site connections amplify the exposure: 61% report limited or no visibility into traffic flowing across permanent tunnel links between offices, data centers, and partner networks. 41% still route vendor and contractor connections through this same infrastructure, typically granting the same access as employees. Every one of these connections extends the blast radius of a single breach.

In contrast, an access model that connects each session only to its intended application leaves a compromised credential with nowhere to go.

If a threat actor successfully compromised a single remote user's VPN credentials today, how much of your network would they have access to?

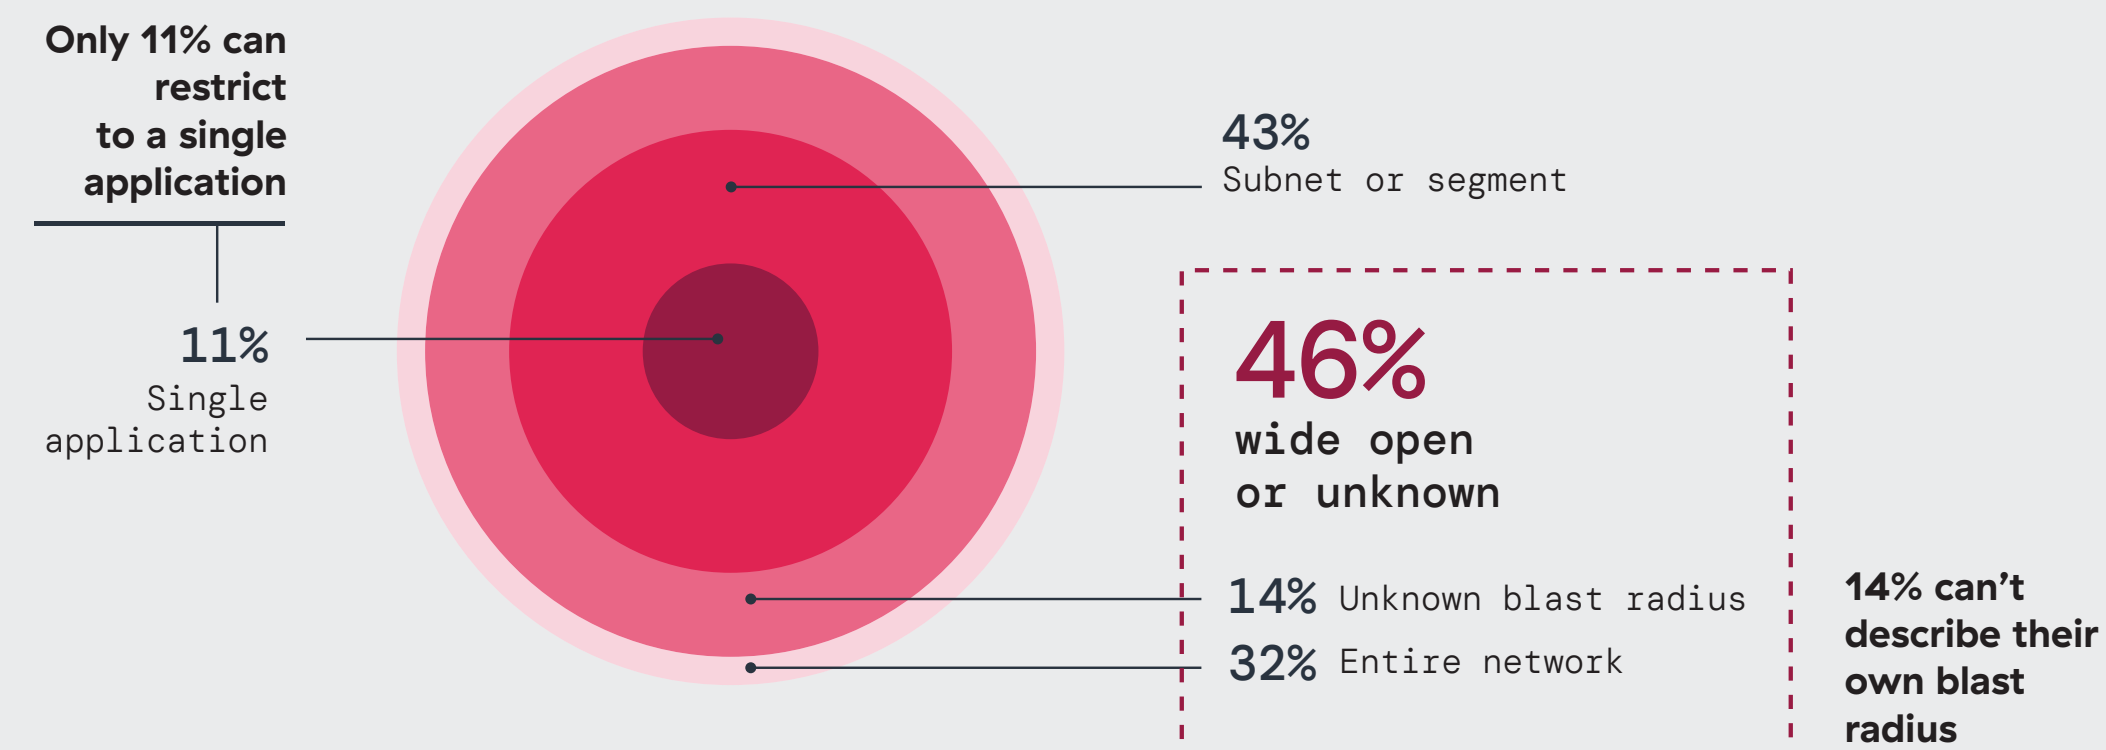


Figure 7: One Credential, Wide-Open Access

Does your VPN infrastructure enable app-level segmentation for different user groups?




Figure 8: Segmentation Stops at the Network Layer



One Credential, Full Access

Composite — Attack risk, Architecture concern, Ransomware vector, Concern level

#1 ACROSS EVERY RISK DIMENSION

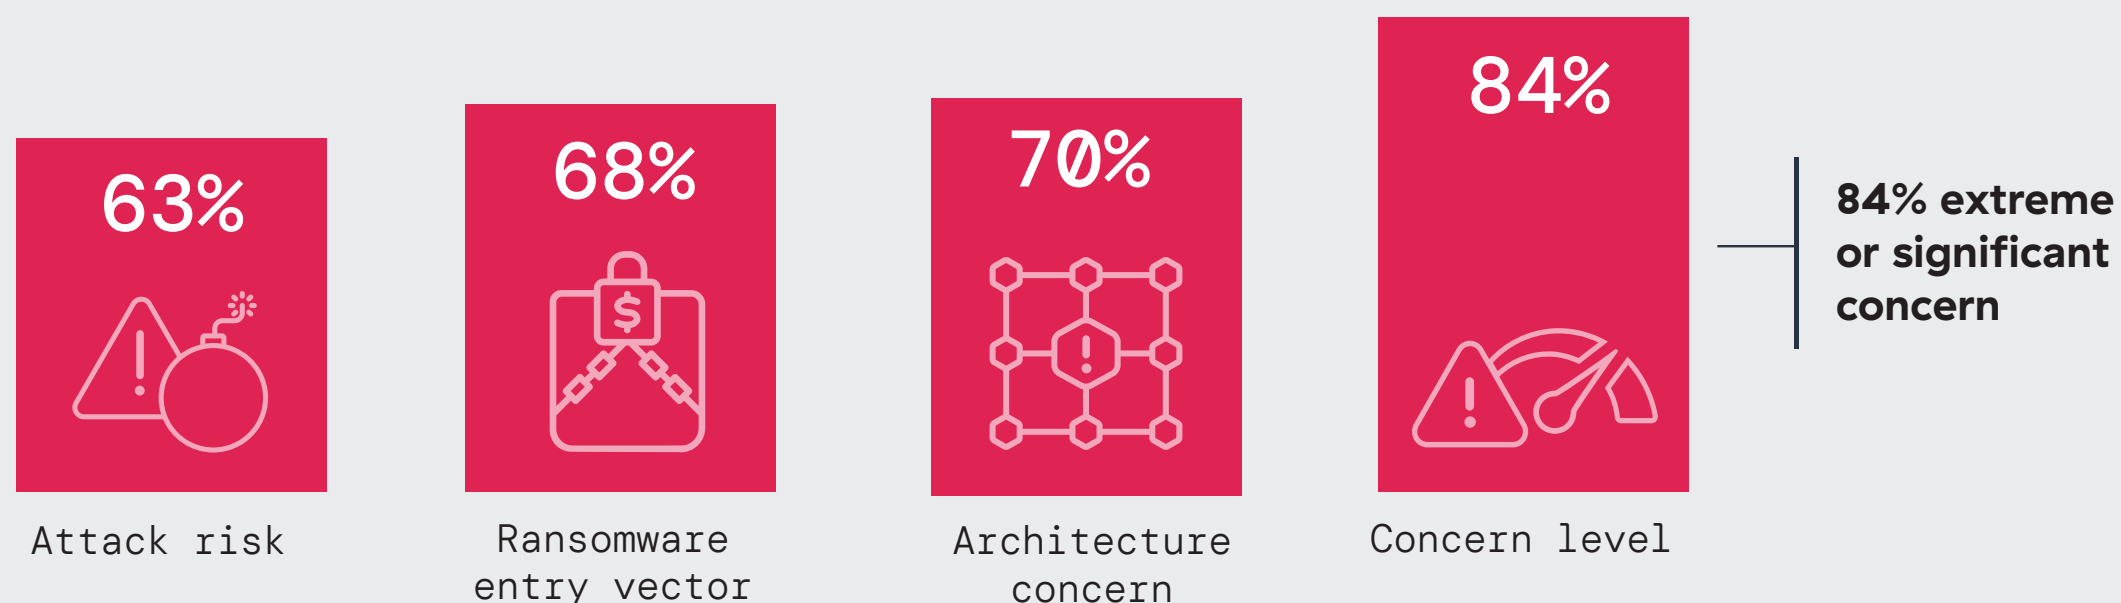


Figure 9: Credentials: #1 Across Every Risk Dimension
Multiple responses allowed

Does your organization have processes in place to regularly audit application access rights for VPN users (including third parties)?




Figure 10: Third-Party Access: Unaudited, Unmanaged

In most breach investigations, the root cause lands in the same place: a stolen credential that opened the front door. VPN makes that credential uniquely dangerous.

Credential theft dominates every risk dimension the survey measured. As the greatest VPN attack risk, 63% named it first. As an architecture concern, 70% ranked it the top threat. As a ransomware entry vector, 68% placed it above all alternatives. 84% expressed extreme or significant concern. No other finding achieved that consistency.

The reason is structural. VPN authenticates once at the perimeter and extends network access from that single event. Anyone holding a valid credential inherits whatever access the legitimate user had, often far more than the role requires. Yet 57% of organizations do not audit the third-party access those credentials unlock.

AI has industrialized the supply side. Infostealer malware now harvests VPN credentials at scale, feeding access broker marketplaces where verified logins are sold with bulk pricing and customer service for ransomware affiliates. In December 2025, GreyNoise detected a coordinated brute-force campaign targeting Cisco SSL VPN and Palo Alto GlobalProtect portals, with over 10,000 unique IPs probing exposed authentication endpoints in a single week.⁶

Dormant vendor accounts, provisioned months ago, never audited, still active, rank among the most valuable listings on access broker forums. Every unaudited credential is an open invitation that never expires.

Continuous identity verification – evaluating device posture and behavior on every request, not just at the perimeter – transforms authentication from a single event into a persistent control.

6. GreyNoise Intelligence, Mass Exploitation Campaign Targeting VPN Portals, December 2025.



Ransomware's Favorite Door

A ransomware operator needs three things: a way in, room to move, and time to work. VPN provides all three.

The credential exposure and patching gaps documented in this report converge here. Once inside, the attacker lands on a segment with access exceeding what the role requires. Without meaningful segmentation, the pivot from entry to encryption is unobstructed.

The Fog ransomware group completed full chains from VPN login to encryption in under two hours.⁷ A joint FBI and CISA advisory confirmed Akira collected over \$42 million from VPN exploitation campaigns; the group surged again in mid-2025 targeting SonicWall SSL VPN devices.² Black Basta's leaked internal communications read like procurement documentation: vulnerabilities ranked by exploitability, credential pricing from access brokers, target lists filtered by internet-facing portals.⁸

Multiple campaigns in 2024 and 2025 demonstrated that implanted malware can survive reboots and firmware upgrades on VPN appliances, meaning a compromised concentrator maintains access indefinitely through a tunnel designed never to close.

The survey data maps onto this kill chain precisely: broad access in 56% of environments, flat topologies in 32%, zero inspection in a third. 83% rank ransomware through VPN as a top-tier concern — yet 53% report no meaningful segmentation to slow it down. The concern is nearly universal; the capability to act on it barely exists.



Until every stage of that kill chain meets a control that stops it — continuous identity verification, inline traffic inspection, per-session application access, and automated threat detection — the path from first credential to final payload stays clear.

In your opinion, what is the primary weak point or entry vector for ransomware in your VPN environment?




Figure 11: How Ransomware Gets In Through VPN
Multiple responses allowed

Under 2 hours from login to encryption

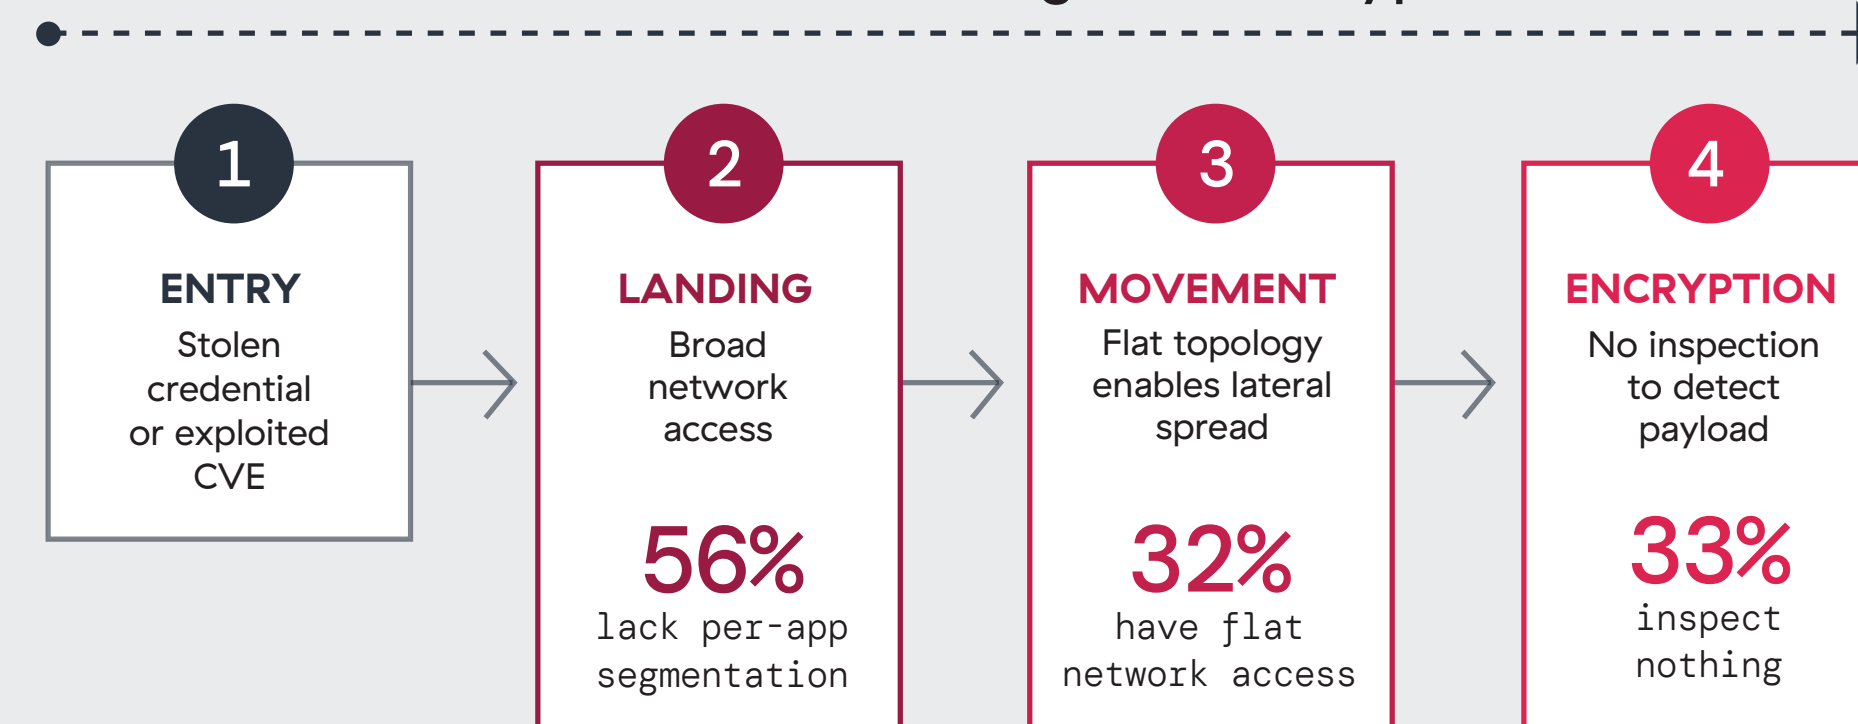


Figure 12: The VPN Ransomware Kill Chain

7. Arctic Wolf Labs, Fog Ransomware Campaign Analysis, 2025. Full chain from VPN login to encryption in under two hours.

2. FBI and CISA, Joint Cybersecurity Advisory: Akira Ransomware, April 2024. \$42M+ in proceeds from VPN-based campaigns.

8. Black Basta internal communications leaked in February 2025; analyzed by multiple threat intelligence vendors.

The Operational Tax

Beyond the security risks, VPN imposes an operational cost that adds up quietly in help desk tickets, user frustration, patching cycles, and the workarounds that emerge when the sanctioned access path becomes the slowest one available.

63% report that users intentionally bypass VPN controls to reach applications faster. When nearly two-thirds of the workforce routes around the security perimeter, every bypassed session flows uninspected and unmonitored. Controls that users consistently route around do not meaningfully reduce risk.

The bypass has roots in measurable friction. Slow connections top the complaint list at 30%, followed by inconsistent device behavior at 23% and frequent disconnections at 20%. All three are performance problems, accounting for 73% of complaints. Forty-five percent report significant or major productivity impact. These are real costs: delayed projects, missed collaboration windows, frustrated employees working around their own security tools.

Security teams spend more time maintaining VPN than defending what sits behind it. 73% say VPN demands more operational effort than modern alternatives. The top drivers: patching complexity (56%) and integration struggles with modern security stacks (50%). As one survey respondent put it, **more time goes to patching the entry point than securing the actual data behind it.**

VPN now creates more exposure through workarounds than it prevents through protection. That is no longer a security problem; it is an operational failure.

What is the most common complaint your users have regarding VPN access today?

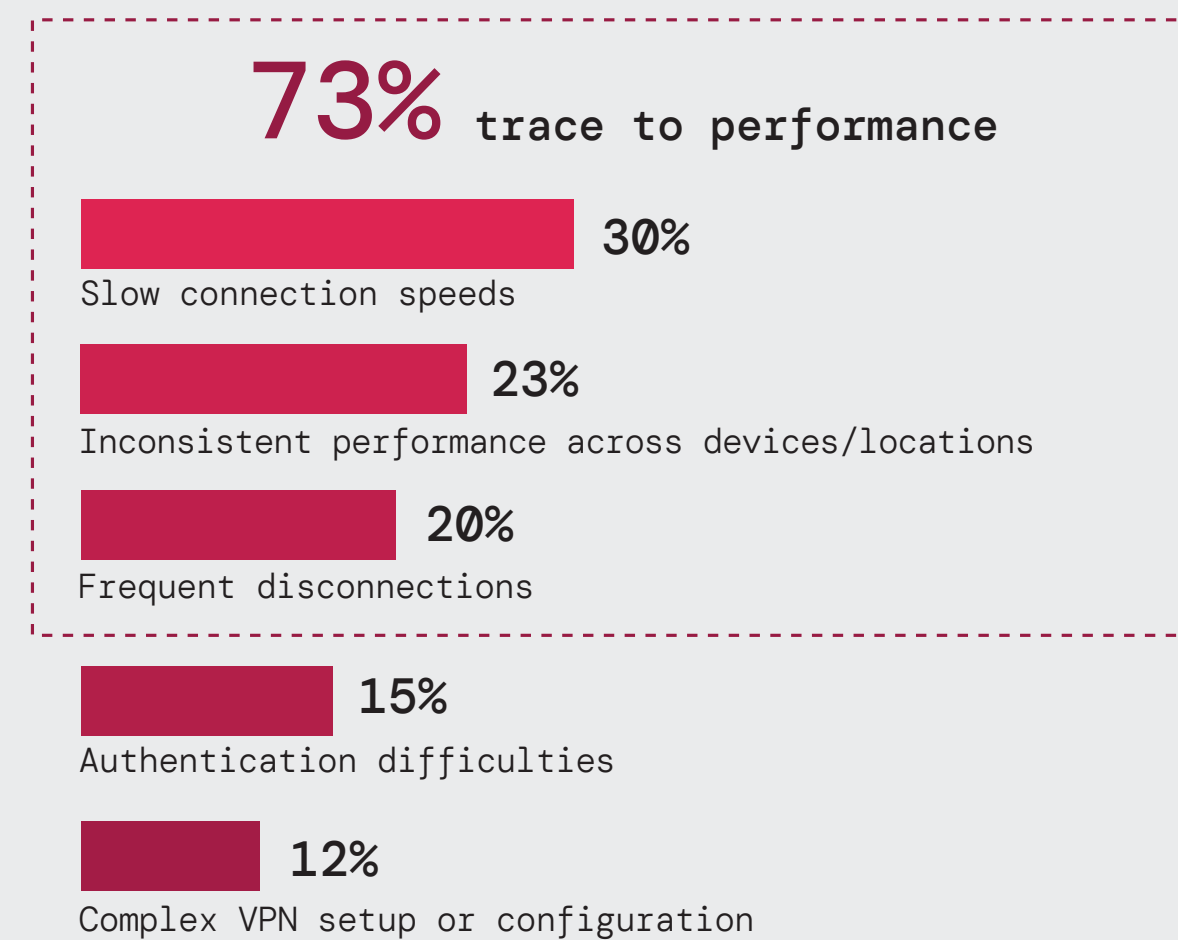


Figure 13: What Users Complain About Most

How frequently do users attempt to bypass or work around VPN access controls to access applications faster?




Figure 14: 63% of Users Route Around VPN



Nation-States Inside the Perimeter

VPN has become a national security concern. The campaigns disclosed in 2024 and 2025 made that clear.

Volt Typhoon, confirmed by CISA and multiple intelligence agencies, targeted critical infrastructure across communications, energy, transportation, and water sectors, maintaining access within some networks for years before discovery.⁹ Salt Typhoon achieved deep penetration of U.S. telecommunications providers, compromising the lawful intercept systems used by law enforcement. The FBI characterized it as one of the largest intelligence compromises in American history.¹⁰

These were not isolated incidents. A China-aligned group designated UAT4356 exploited three Cisco ASA/FTD zero-days from May through September 2025, targeting government networks worldwide with firmware-level persistence that survived reboots and upgrades.¹¹ Additional Chinese state-sponsored groups exploited Ivanti VPN zero-days throughout early 2025, deploying malware designed to persist despite standard remediation.

The CrowdStrike 2026 Global Threat Report found that 40% of vulnerabilities exploited by China-nexus actors targeted edge devices such as VPNs, firewalls, and gateways.¹


In this survey, 51% of organizations reported a VPN-related security incident in the past twelve months — a figure that spans every sector, not just the critical infrastructure these campaigns targeted.

CISA issued consecutive emergency directives in January 2024 and January 2025, ordering federal agencies to disconnect affected appliances after malware proved capable of surviving factory resets and firmware updates.¹²

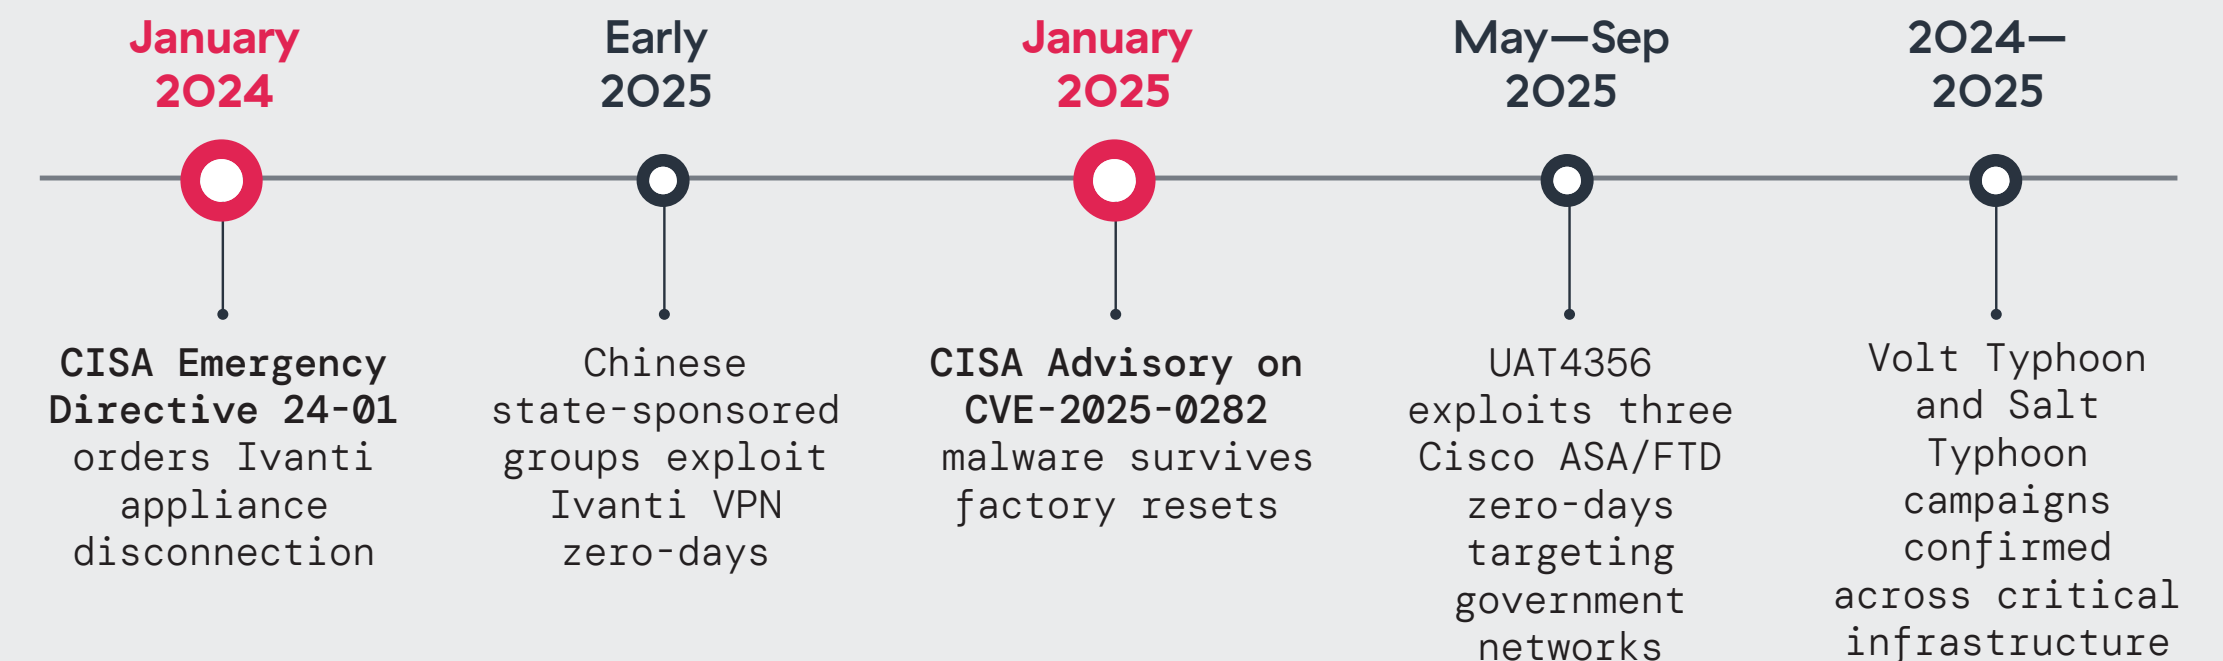
When standard recovery fails and the government orders disconnection, the security device itself has become the threat — ordered off the network it was purchased to protect.

Against adversaries with nation-state capability, the only viable defense is an architecture that never exposes the appliance in the first place.

Has your organization experienced a security incident related to VPN vulnerabilities in the past 12 months?



The Cadence of Nation-State VPN Exploitation



9. CISA, Advisory on Volt Typhoon, multiple releases 2024–2025; confirmed by FBI, NSA, and Five Eyes intelligence partners
 10. FBI Director Christopher Wray, Congressional testimony on Salt Typhoon telecommunications compromise, 2024.
 11. Cisco Talos, ArcaneDoor Campaign Disclosure (UAT4356), September 2025. Three ASA/FTD zero-days exploited from May–September 2025.
 1. CrowdStrike, 2026 Global Threat Report, February 2026.
 12. CISA, Emergency Directive 24–01 (January 2024) and CISA Advisory on CVE–2025–0282 (January 2025).



The Clock Is the Risk

Every finding in this report reduces to one variable: time. Time to patch, time to detect, time to contain.

Only 6% can patch a critical VPN vulnerability within 24 hours, while 79% say their top AI-driven risk is attackers weaponizing vulnerabilities faster than patches can be deployed. The exposure window is structural, not temporary.

With 38% transitioning and 34% planning zero trust, hybrid access will persist for years. During that coexistence period, broad network exposure remains in production.


The priority is speed of containment: shrinking reachable surface area, enforcing per-application access, and restoring inspection at the point of connection.

A CISA-aligned readiness assessment quantifies how much exposure remains. The decisive question is whether your access architecture can contain the next compromise before it spreads.



6%

can patch
within
24 hours



29_{min}

average
breakout
time



From Reactive To Resilient

Most organizations in our survey placed themselves in the Reactive column across every dimension. That position is not sustainable when breakout times are measured in minutes and AI-enabled attacks are reported by 61% of respondents.

These four changes define the shortest path from Reactive to Resilient. Each maps to a CISA zero trust pillar and is ordered by impact.

1. First: Eliminate broad network access (Networks):

32% still grant flat network access after authentication, and 24% provide effectively open connectivity. Shifting to an access model that connects each session only to its intended application is both an architecture change and a segmentation fix in a single step: it collapses blast radius for credential compromise and sharply reduces lateral movement pathways. Least-privilege access is foundational, not a later phase. It directly addresses the risk that 67% cite as their primary reason for adopting zero trust.

2. Second: Make authentication continuous (Identity):

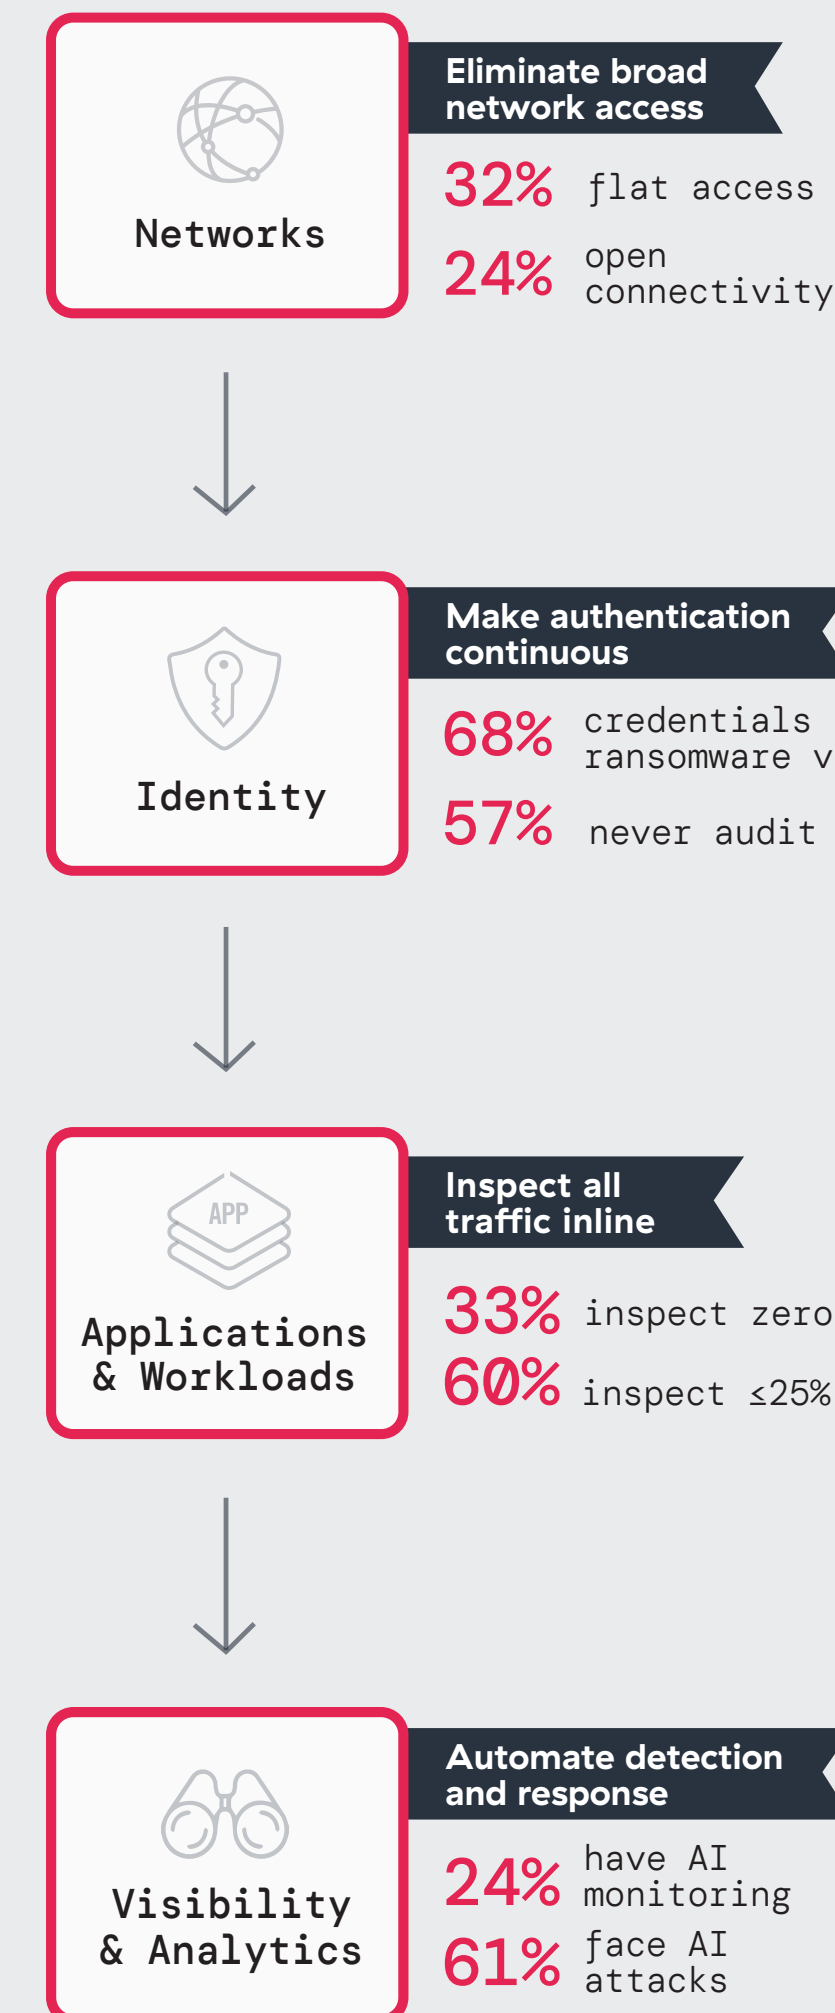
68% identify credentials as the top ransomware entry vector, yet most VPN architectures validate identity once and never again. Fifty-seven percent do not audit third-party VPN access or are unsure whether audits occur. Continuous verification that evaluates identity, device posture, and behavior on every request transforms authentication from a single checkpoint into a persistent control — closing the gap that credential-based attacks exploit.

3. Third: Inspect all traffic inline (Applications & Workloads):

One-third inspect zero encrypted VPN traffic, and 60% inspect a quarter or less. Without inspection, every session is a potential channel for malware, C2, and exfiltration. Inline inspection at the access layer restores the visibility that encrypted tunnels strip away and makes applications invisible to unauthorized discovery.

4. Fourth: Automate detection and response (Visibility & Analytics):

24% deploy AI-powered monitoring today, yet 61% have encountered AI-enabled attacks. Over half say legacy VPN infrastructure blocks integration of AI-driven security tools entirely. An access model built for automated detection and response removes the integration barriers that keep most organizations blind to the fastest-moving threats.



None of these changes require perfection across every dimension at once. Each one independently narrows the gap an attacker can exploit, and every week of delay keeps the full gap open.

Is your organization actively transitioning away from VPNs toward a Zero Trust security model?

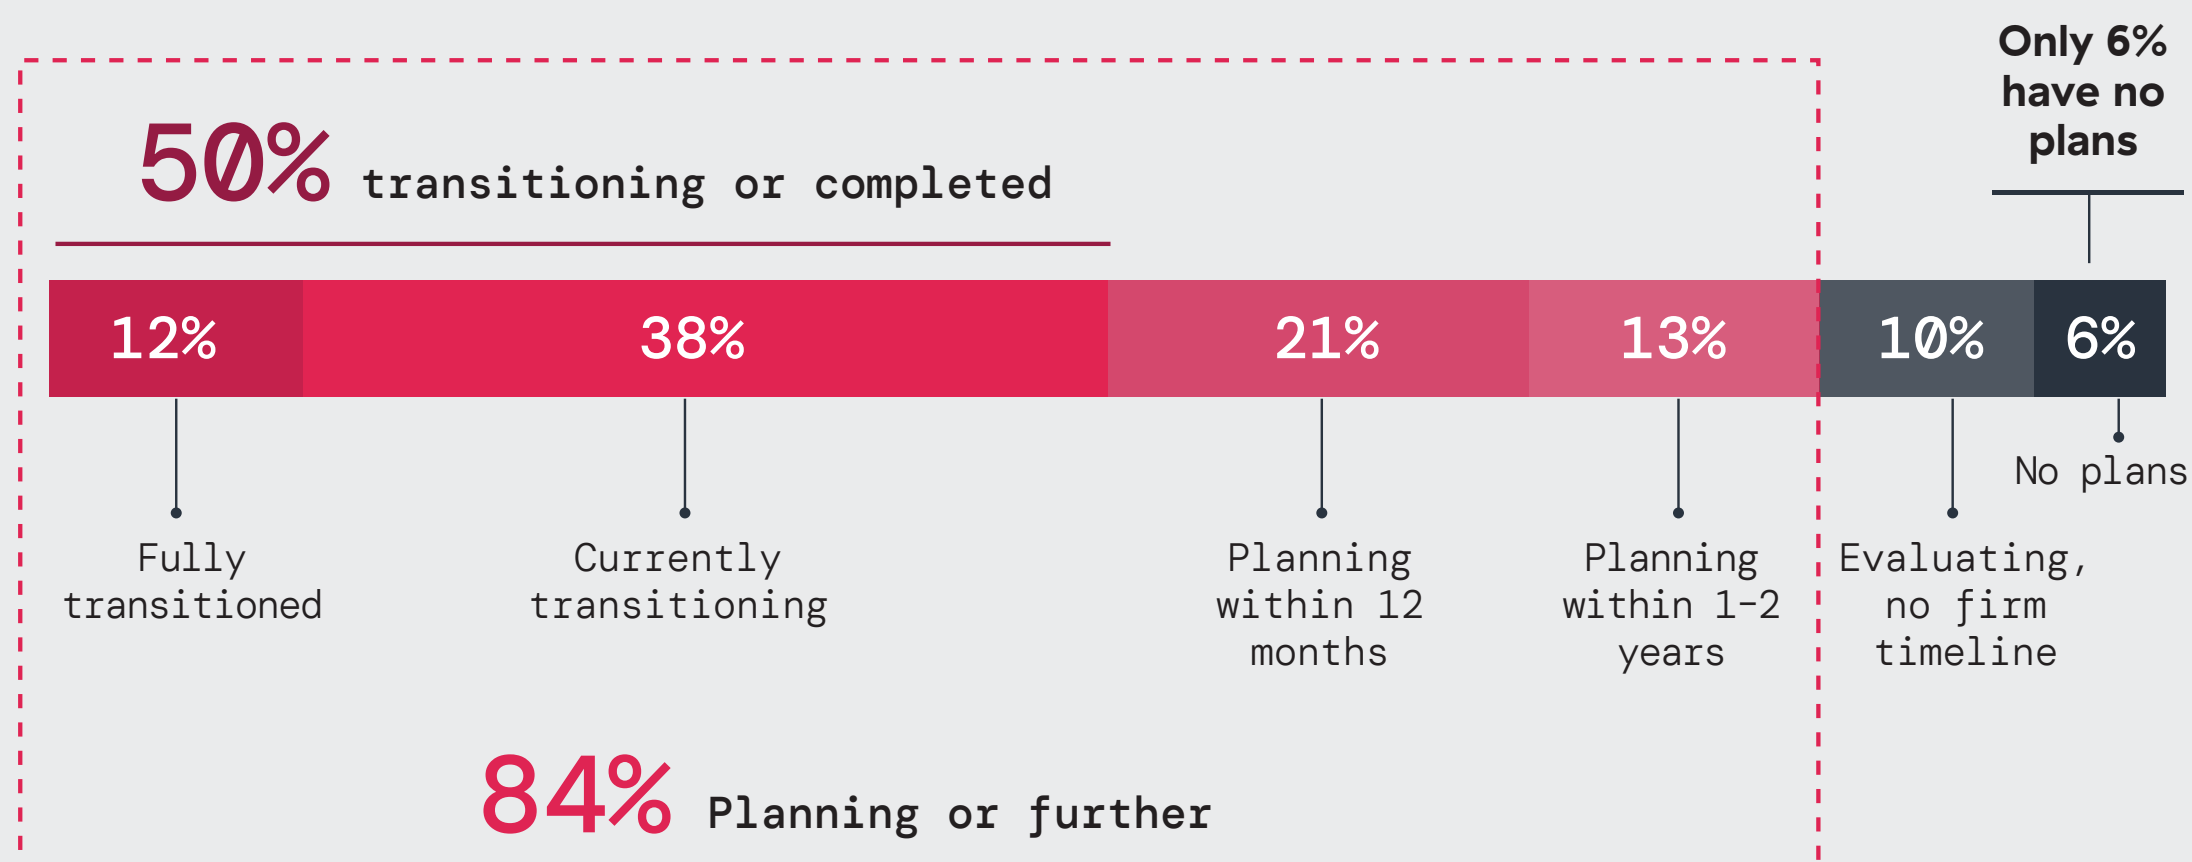



Figure 16: 84% Are Moving Beyond VPN

Steady acceleration across three years



Three consecutive surveys

The contrast is architectural. VPN authenticates once and exposes the network. Zero trust verifies continuously and connects each session only to a single application, keeping the internal network unreachable. At its core, it's the shift from putting users on the network to connecting them only to the specific app they need.

Zero Trust The Execution Gap

The strategic debate over whether to move beyond VPN is effectively settled. The hard part now is execution: how fast organizations can transition, and how much exposure accumulates during coexistence.

VPN usage is declining in 51% of organizations, and only 10% report an increase. Half are actively transitioning to zero trust or report completion, and another 34% are in formal planning. Across three consecutive surveys, the share at least planning zero trust has grown from 78% to 81% to 84%. Only 6% report no plans.

The drivers map directly to this report's findings. Eliminating lateral movement leads at 67% — the same threat that 81% identified as a top concern and 77% cannot contain. Operational overhead reduction follows at 62%.

Better segmentation at 58% and ransomware protection at 53% complete the picture. The alignment between what organizations fear and what drives their migration to zero trust is nearly exact.

Zero trust is a journey, and most enterprises will operate hybrid access environments for years as they migrate high-risk workforce access first while legacy VPN use cases persist. The practical objective is to reduce exposure in measurable steps by shrinking reachable network surface area and limiting blast radius during coexistence.



VPN to Zero Trust: Readiness Assessment





AI has compressed attack timelines across every dimension of enterprise security. Credential harvesting, exploitation, east-west movement, and exfiltration all operate faster than just twelve months ago. The readiness question is whether your infrastructure can operate at the speed these threats now demand.

This assessment is structured around the CISA Zero Trust Maturity Model v2.0, scoped to the access risks this survey measured. Two CISA pillars, Devices and Data, fall beyond this survey's scope; organizations should assess those independently. The three maturity stages align to the progression CISA describes from Traditional through Optimal. In every dimension, the majority of respondents fall in the Reactive tier. The lowest-rated dimension typically represents the fastest path an attacker will take.

For each dimension, identify which column most closely describes your current operational state — not your planned state or your policy, but what is enforced in production today. Many organizations will find themselves in the Reactive column across multiple dimensions. That honest placement is the starting point for prioritization.


The next page maps the four changes that move each dimension from Reactive to Resilient, ordered by impact.

Zero Trust Readiness Matrix (aligned to CISA Zero Trust Maturity Model v2.0)

DIMENSION (CISA Pillar)	REACTIVE	TRANSITIONING	RESILIENT
 Identity	THE MAJORITY IS HERE ▼ Single login grants persistent access. 68% cite credentials as top ransomware vector; 57% never audit third-party access	MFA and periodic reauthentication for some applications; third-party access audited annually	Continuous verification on every request; adapts to identity, device posture, and behavior
 Networks	Broad network access post-auth; flat topology. 56% lack per-app controls; 77% cannot contain lateral movement	Zone or role-based segmentation; sessions not yet isolated	Each session connects to one application only; lateral movement eliminated by default
 Applications & Workloads	Little or no encrypted traffic inspection; apps reachable by anyone on-network. 60% inspect ≤25% of traffic; 70% blind to AI threats in VPN	Some encrypted inspection; manual anomaly review	All traffic inspected inline in real time; apps invisible to unauthorized users
 Visibility & Analytics	No automated detection; legacy infra blocks modern tools. 76% lack AI monitoring; 51% blocked by legacy VPN	Known-pattern detection; manual investigation and response	Automated detection and response; telemetry feeds continuous policy refinement

The lowest-rated dimension typically represents the fastest path an attacker will take.

Predictions For 2026 And Beyond



1

Credential theft becomes a default VPN intrusion path.

AI-generated phishing/social engineering will outpace exploitation because “log in” is faster, quieter, and more reliable than “break in.”

2

Attackers optimize for machine-speed time-to-compromise.

AI will continuously refine exploit targeting and follow-on actions to compress the window from first contact to valid VPN access—from days to hours, increasingly minutes.

3

MFA, resets, and help desk flows become the new frontline.

AI-assisted social engineering will target push fatigue, one-time codes, device enrollment, and password reset workflows as aggressively as passwords.

4


Patch velocity won't match exploitation velocity.

AI-assisted scanning and exploitation will shorten the time between disclosure and compromise attempts, widening the gap for organizations with slower patch cycles.

5

Defenders shift to continuous verification—and away from VPN.

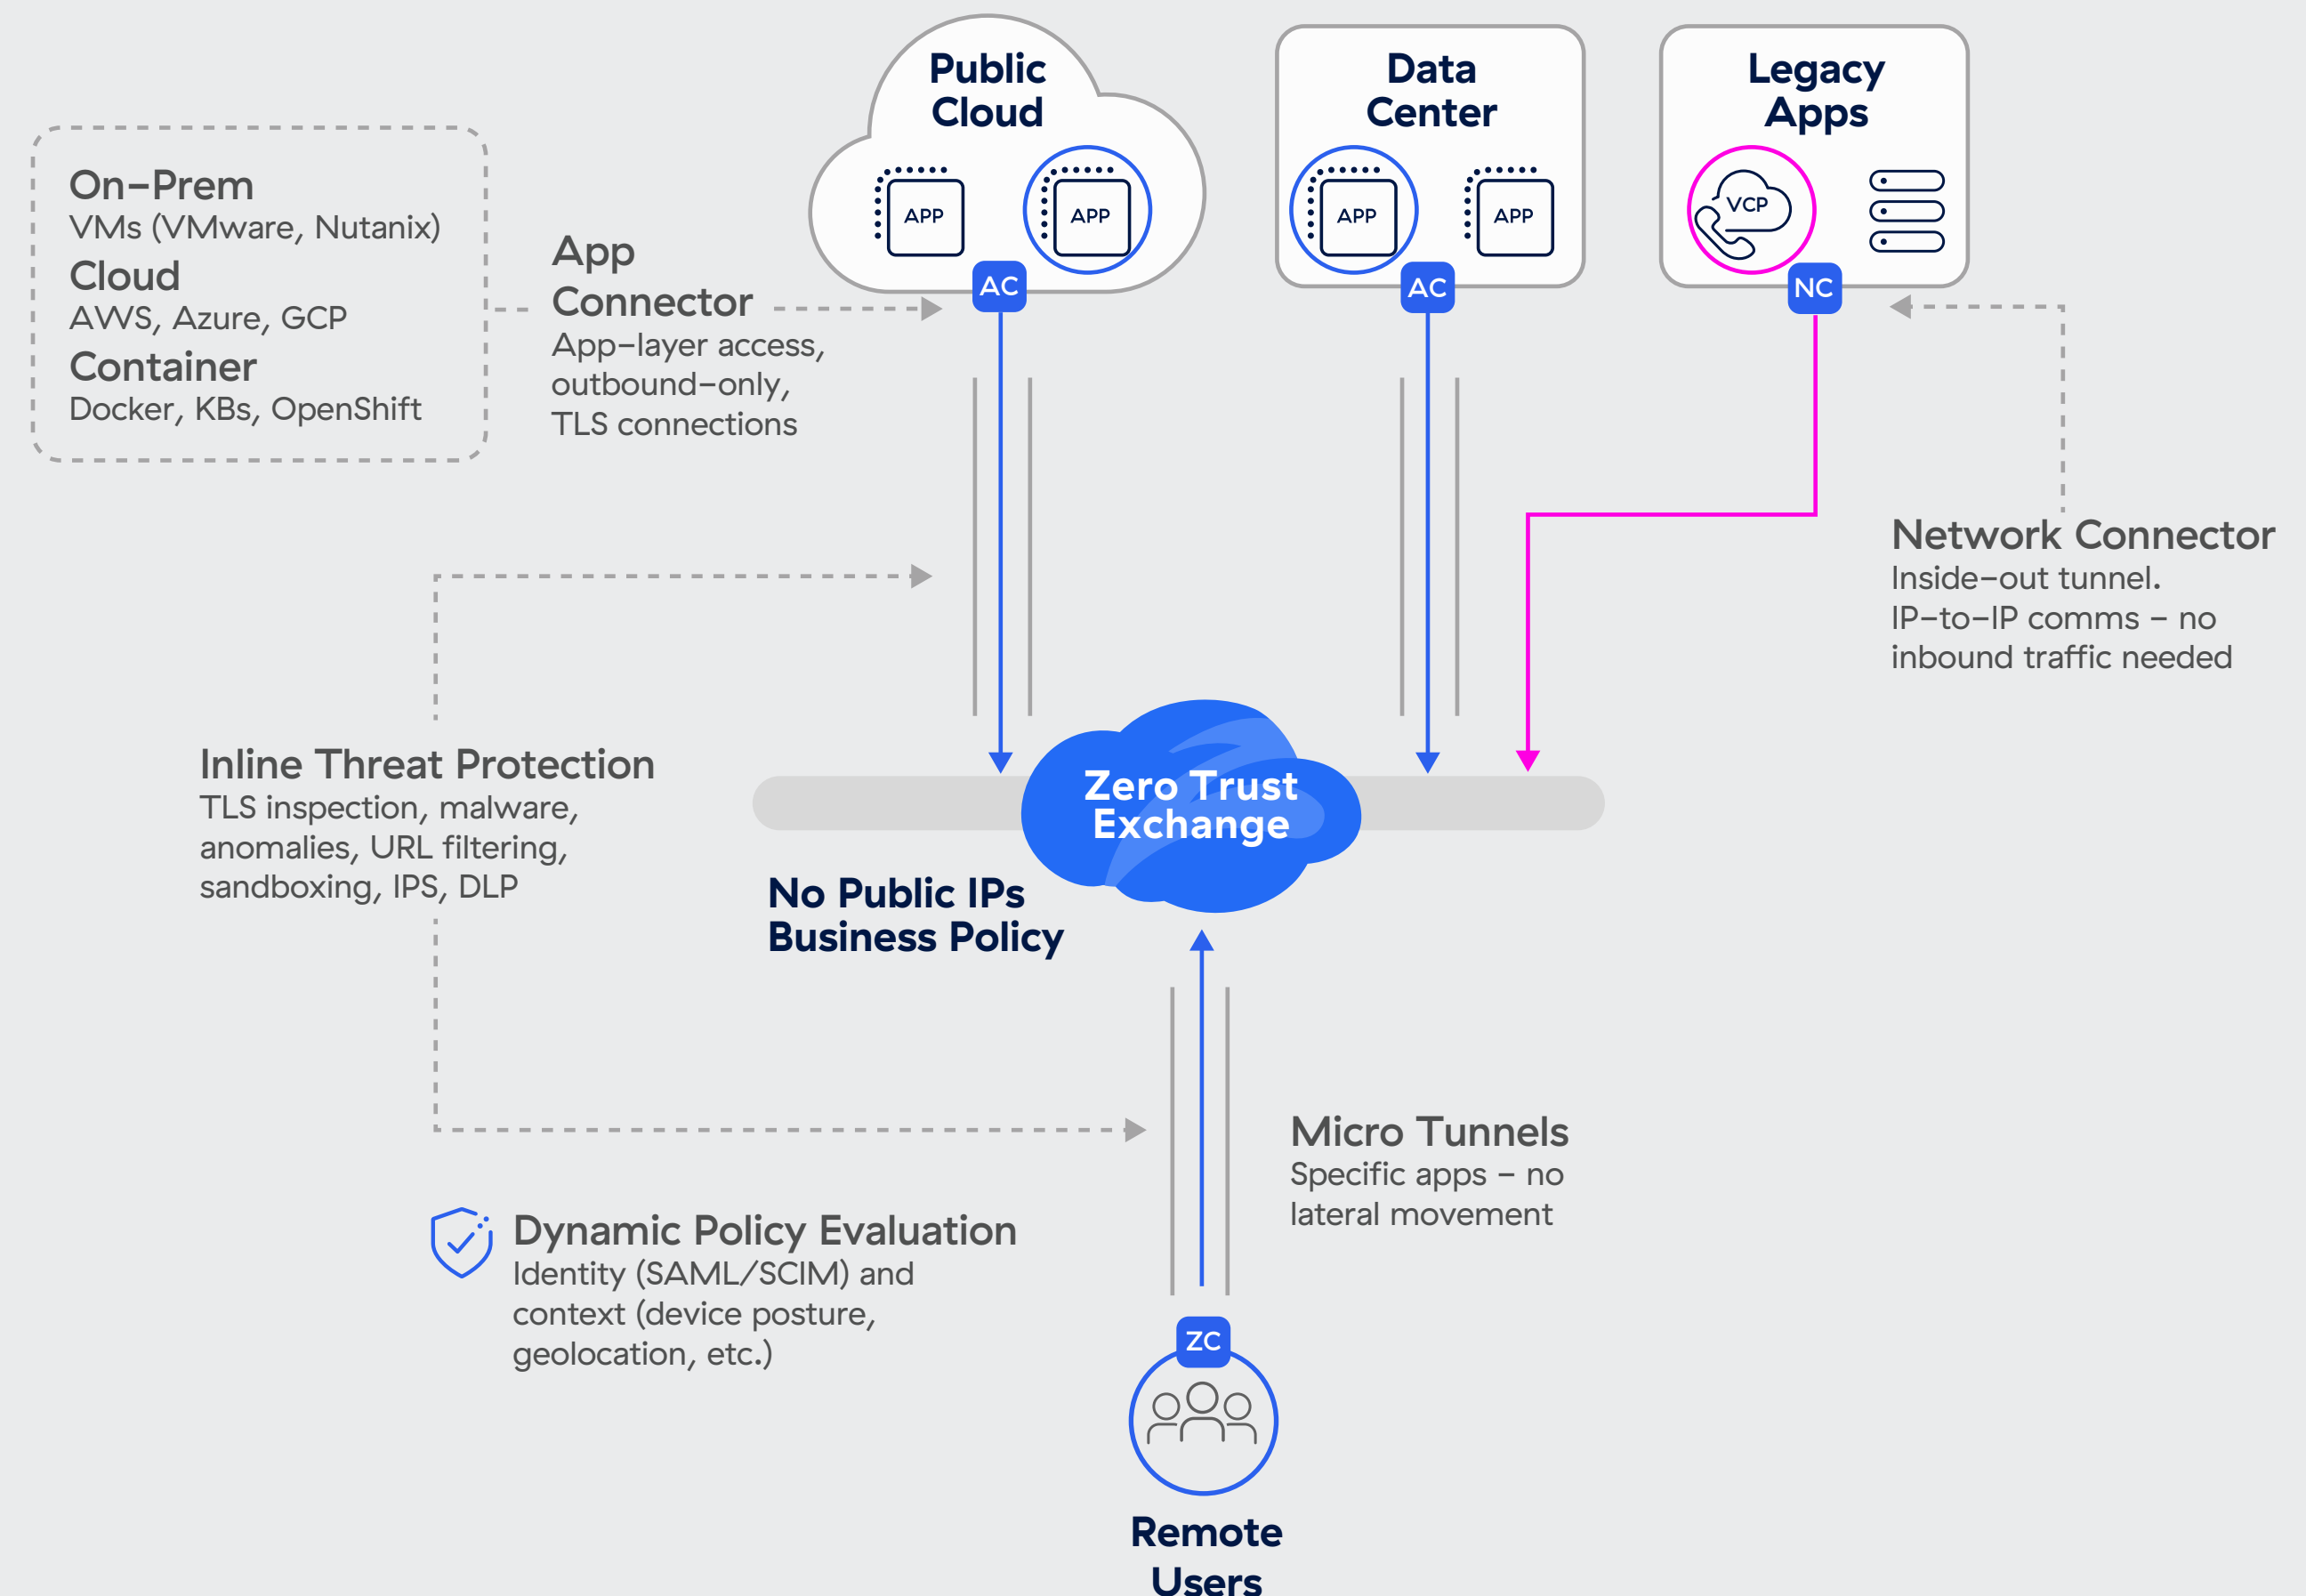
Organizations will move from “authenticate once” perimeter trust to per-request checks (identity, device posture, behavior) and least-privileged user-to-app access to survive machine-speed attacks.



How Zscaler Transforms Secure Access

If time is the risk, then secure access has to reduce time everywhere it matters: time to expose services, time to move laterally, and time to contain. Legacy VPN does the opposite. It extends the network to the user, broadens the attack surface, and turns a single credential compromise into a pathway for lateral movement. Zscaler Private Access (ZPA), the industry's first AI-powered ZTNA, replaces network-level access with direct user-to-application connectivity. Applications are made invisible behind the Zero Trust Exchange, and Autonomous User-to-App Segmentation enforces least-privilege access by default, collapsing blast radius and helping eliminate lateral movement.

ZPA supports the broadest set of private access use cases from managed and unmanaged (BYOD) devices to in-office users, third-party contractors, and privileged access (RDP, SSH, VNC) without reintroducing a flat network. Built to deliver Zero Trust Everywhere at scale, it provides consistent security and performance across branch, home, and cloud on a unified platform for users, workloads, and OT devices, backed by 160+ global data centers. And because fast attacks require fast containment, ZPA adds inline inspection plus integrated AppProtection, deception, and rapid deployment and scaling so organizations can reduce exposure and improve resilience during the hybrid transition years.





Key Differentiators Of Zscaler Private Access (ZPA)

- **Built from the ground up for least-privileged access**
Connect users to approved apps—not the network—with policy-based access by default.
- **Minimize the attack surface and lateral movement**
Keep apps invisible and contain risk with user-to-app segmentation.
- **Full inline inspection**
Block threats and protect data with inline inspection and integrated DLP.
- **Universal ZTNA**
Deliver consistent access to any app (including legacy) from anywhere.
- **Boost workforce productivity**
Provide fast, direct access—no backhauling through the data center.
- **Secure B2B connectivity**
Give partners on-demand, zero-trust access—without added infrastructure.
- **Accelerate M&A/Divestiture**
Integrate or separate access quickly—without merging networks.
- **Reduce operational complexity**
Replace VPN/VDI/firewalls with a cloud-native access service.
- **Automated business continuity**
Maintain policy-enforced access through outages and disruptions.
- **Centralized policy enforcement**
Apply one policy model to all apps for remote and in-office users.
- **Cloud-delivered solution**
Connect users directly to apps without routing through the data center.
- **160+ global cloud edge locations**
Deliver secure, local access with consistent performance worldwide.

Why Zscaler Private Access (ZPA)?

Minimize Attack Surface

Applications are invisible to the network making them impossible to discover or attack.

Boost Productivity

Fast, direct access from 160+ PoPs worldwide with no traffic backhauling.


Reduce Total Cost of Ownership (TCO)

Consolidate VPNs, firewalls, and load balancers into a single cloud-native platform.




Methodology & Demographics


Career Level




Department



Company Size



Industry



This report is based on a comprehensive survey of 822 IT and cybersecurity professionals conducted in early 2025 by Cybersecurity Insiders in partnership with Zscaler ThreatLabz. The research examines how organizations are addressing the intersection between AI-accelerated threats and traditional perimeter-based access, encompassing VPN vulnerability exposure, lateral movement risk, and readiness for zero trust migration. Respondents were screened for direct operational or strategic responsibility over their organization’s network access architecture. Using a stratified sampling approach, the survey achieved a 95% confidence level with a margin of error of +/- 3.4%.

Note on terminology: Throughout this report, “AI-enabled attack” refers to adversary activity where responders observed either automation-driven speed/scale consistent with AI augmentation or direct use of AI-generated content in social engineering. Confirmed and suspected incidents are reported separately.

©2026 Cybersecurity Insiders. All rights reserved.

Limited editorial citation (up to 100 words and one unaltered chart) is permitted with clear attribution to “Cybersecurity Insiders, 2026 VPN Risk Report” and a visible link to cybersecurity-insiders.com.

The report sponsor may reference the findings and use individual charts or data points in presentations and marketing materials with proper attribution. The full report, underlying dataset, and research methodology remain the intellectual property of Cybersecurity Insiders and may not be reproduced, redistributed, or incorporated into derivative research without written permission.

This report was produced by Cybersecurity Insiders with the support of Zscaler. Permissions: info@cybersecurity-insiders.com



Cybersecurity

INSIDERS

BENCHMARK YOUR SECURITY MATURITY

Independent cybersecurity research revealing the gaps that shape cybersecurity strategy

Cybersecurity Insiders produces independent research based on surveys of cybersecurity leaders and practitioners worldwide. Our reports reveal where security strategies break down in practice — helping organizations benchmark their maturity, identify capability gaps, and prioritize the actions needed to close them.

For more information, visit

cybersecurity-insiders.com



Zero Trust Everywhere

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange™ platform protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 160 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at zscaler.com or follow us on X @zscaler.

© 2026 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at zscaler.com/legal/trademarks are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.

+1 408.533.0288

Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134

zscaler.com