



ThreatLabz 2026: Frontier AI Readiness Report



Table of Contents



Executive Summary	3
By the Numbers	5
Assessment Cohort	6
Threat Outlook	7
Methodology	9
Key Findings	10
The New Attack Chain: How Frontier AI Autonomously Chains Small Gaps Into a Single Exploit Sequence	15
Case Evidence: Three Organizations in Progress	16
Implementation Realism: What the 90-Day Sprint Actually Requires	18
Recommendations	19

Executive_Summary

Last year, we received a call from one of the well-known AI companies. We had just been notified about their release of a frontier AI model tuned to execute autonomous security testing. That call set us on a journey to help organizations prepare by examining their readiness against these autonomous attacks. We were just as surprised by the findings as we were by its capabilities.

Of note, we didn't apply Frontier AI to organizations' environments; instead, we first tested it on our own infrastructure, using those insights to develop the assessment program we now offer.

The goal is not just to call out gaps, but to help illustrate where security can be strengthened as organizations prepare for a new wave of threats. We are all learning together, and these findings are meant to serve not just as a warning, but as a shared lesson in how we can improve resilience across the board.

Because Zscaler operates a cloud-native architecture unlike the hybrid and legacy environments most enterprises run, we used the same frontier-testing methods seen with real models and replicated them with similarly effective models and TTPs to emulate how these adversaries operate in organizations' environments. This allowed us to build and calibrate the assessment program for the operational realities of enterprise infrastructure, rather than for architectures like our own.

Our biggest finding: The winning response to AI is not simply more AI, but architecture that denies exploitable paths by design.

Global enterprises are not prepared for frontier AI-powered cyberattacks — but not for the reason most people assume. The danger isn't AI uncovering exotic, never-before-seen vulnerabilities. It's that frontier models can take the obvious ones — the gaps every organization already knows exist — and chain them into complete attack paths at machine speed. Most architectures were never designed for that, and trying to outrun AI — whether with humans or tools — is an ever-losing battle. Architecture is the only thing that wins. The good news is that with the right zero trust architecture in place, most enterprises are closer to resilience than they think.

The Frontier-AI-Readiness Assessment Program measures organizational resilience against a specific kind of adversary: a frontier AI model operating at machine speed and capable of autonomous reconnaissance, vulnerability chaining, and lateral movement. We use the term "Frontier-AI" to describe this adversary class because it can reason across attack surfaces in ways that traditional scanners and even human analysts cannot replicate at comparable speed. In parallel testing against production environments, frontier models showed high signal-to-noise reasoning with grounded context, effectively serving as force multipliers for attack chain discovery. These capabilities are not theoretical; they have already been demonstrated.

Across hundreds of enterprise organizations assessed in H1 2026, including Fortune 500 companies in financial services, healthcare, manufacturing, energy, and eight additional industries, the average Frontier-AI-Readiness Score was 37 out of 100. Almost all organizations (91%) had publicly reachable AI infrastructure, and none of them (100%) had identity-based access controls in place. Additionally, not a single organization inspected enough encrypted traffic for its security stack to function as designed. The average enterprise was effectively blind to 64% of the data traversing its own network, and none had closed the gap between owning defensive capabilities and actually operating them.

The most urgent emerging gap compounds the structural finding above: AI governance is at year zero across every assessed enterprise. AI surfaces, including copilots, inference APIs, MCP endpoints, and shadow AI gateways, are proliferating faster than security programs can absorb and govern them.

This presents a significant opportunity to elevate readiness against frontier AI threats. We found that nearly a third of organizations could dramatically improve protection by extending deception to critical choke points. Similarly, moving intrusion prevention into full enforcement mode, integrating application-layer anomaly detection, and refining segmentation policies represent concrete, high-impact steps to improve SOC alerting fidelity and overall defensive posture.

Security leaders should read this report as a catalyst for board-level discussion. Moving from the cohort average of 37 to a target Frontier-AI-Readiness Score between 73 and 98, the threshold at which all five defensive pillars reach operational effectiveness against frontier-AI-class attacks, is achievable within 90 days. It requires three executive actions: (1) establish frontier-AI resilience as an enterprise mandate with board reporting, (2) secure internal executive alignment and dedicated sprint capacity, and (3) execute a dual-track sprint – activate dormant controls and gate every AI surface simultaneously.

The operational starting point is the same for every organization: start with what you already have. The strategic starting point is governance: no organization can defer AI access security without accepting exploitable exposure today.



DEEPEN DESAI
EVP Cybersecurity,
Zscaler



GEORGE MOSER
Global VP & Head of CXO
in Residence, Zscaler

By the_ Numbers

100% of assessed organizations lack identity-based authentication for their AI.

91% of organizations with publicly reachable AI infrastructure are leaving it completely exposed to outside attack.

37 out of 100 is the average starting Frontier-AI-Readiness Score. The floor was 17. The ceiling was 60.

SSL inspection rates range from **0.4% to 68.1%** with an average of approximately **36%**. This sits well below the 70% inspection level considered necessary for acceptable security coverage with inline security controls (DLP, IPS, sandboxing, AI egress).


483,250 applications were reachable by the most entitled identity in the worst-case blast radius scenario.

In a representative wildcard-specific case, 195 wildcard domains alone exposed more than **108,033** applications. (A wildcard is a symbol such as “*” or a leading “.” used in policies to represent multiple domains or paths at once, introducing administrative efficiency but drastically widening the attack surface when not adequately controlled.)

9 out of 10 organizations have exposed VPN appliances. Each is a pre-authentication attack surface that a frontier AI model can probe at machine speed.


750+ assets carry CISA Known Exploited Vulnerabilities on internet-facing infrastructure, confirmed via active probing, with exploits actively used in the wild at the time of assessment.





Assessment_ Cohort

The findings in this report are drawn from Frontier-AI-Readiness Assessments conducted across hundreds of enterprise organizations during H1 2026. 32% are Fortune 500 companies. The cohort spans 12 industries with the following composition:



Assessed organizations were primarily large enterprises with tens of thousands of internet-discoverable assets. Internet-discoverable assets include any system, application, or endpoint reachable from the public internet without authentication.



Threat Outlook

Frontier AI models have demonstrated the ability to identify and chain enterprise vulnerabilities at machine speed.

The findings above describe the current state: conditions observed across the cohort in H1 2026. The threat outlook below explains why the activation timeline matters. The adversary capabilities that exploit these conditions are maturing faster than enterprise operational cycles typically move. Each trend below directly amplifies one or more assessment findings.

In Zscaler's frontier model testing program, Frontier-AI-class models were evaluated across three modes: black-box with no internal knowledge, white-box with source code, compiled binaries, and static files, and gray-box with partial or full system context. These tests independently confirmed that, when provided grounded context, frontier AI models can reason against enterprise environments with high signal-to-noise and consistently identify multi-step exploit chains that human analysts and traditional scanners missed. Across production architectures, grounded context proved to be a force multiplier for attack path discovery.

This independent validation confirms that the configuration gaps identified in Findings 1 through 6 are exploitable by a frontier AI model operating autonomously, and that present-day frontier models can take advantage of these specific conditions as observed across the assessed cohort.

AI-powered vulnerability discovery will very likely accelerate new exploit availability through 2027

Mean-time-to-exploit has gone negative: vulnerabilities are now exploited before patches can be deployed across enterprise environments, even with automated patching in place. The challenge is structural, as exploitation timelines have compressed faster than enterprise change management can adapt.

For internet-facing assets with KEVs and EPSS scores at or near 1.0, as confirmed across the assessed cohort, the window between disclosure and active exploitation has collapsed to hours. Organizations relying on standard maintenance-window patching will almost certainly face exploitation of known vulnerabilities during the remediation cycle.

AI attack surface expansion will likely outpace governance programs for the next 12 to 18 months

The universal absence of AI governance reflects a structural lag between AI deployment velocity and security program adaptation. Development teams, SaaS vendors, and business units are deploying copilots, inference endpoints, and MCP interfaces without security review. Based on the proliferation rate observed, ungoverned AI surface will likely double before governance reaches operational coverage. Organizations that delay will likely face exploitation within the assessment period.

Lateral movement via identity compromise will very likely become the dominant post-exploitation technique against zero trust environments

The wildcard findings indicate identity compromise is the path of least resistance. A Frontier-AI-class adversary will very likely prioritize credential theft over traditional network exploitation. The payoff is disproportionate: thousands of applications via a single identity. Organizations with never-expiring sessions and no device posture validation are almost certainly first targets.

The encrypted traffic blind spot will likely become the primary exfiltration channel for AI-driven data theft

Frontier AI can encode sensitive data within encrypted API calls, model interactions, or cloud traffic. AI tools generating legitimate-appearing patterns combined with security stacks unable to inspect creates a high-bandwidth, invisible exfiltration channel. This is likely already being exploited and will almost certainly intensify.

Autonomous exploit chaining will likely reduce time from initial access to full compromise in minutes

The convergence of exposed VPN appliances, KEV vulnerabilities, wildcard segmentation, and dormant detection creates an environment where an AI adversary can likely progress from reconnaissance to broad internal access without triggering any response. The data suggests this full chain is feasible today against the majority of the cohort. Autonomous tooling executing end-to-end will likely be available to sophisticated actors within 12 months.



Methodology

Frontier AI Threat Validation

We leveraged OpenAI GPT-5.5 and Anthropic Claude Opus 4.8 under controlled harness conditions with grounded enterprise context. We use “Frontier-AI-class” because the specific model matters less than the capability threshold: autonomous multi-step reasoning, tool use, grounded context retrieval, and self-directed exploitation sequencing. Any model at or above that threshold, current or future, represents the adversary class this assessment measures.

In parallel, Zscaler conducted internal testing frontier model testing using across three modes: black-box (zero internal knowledge), white-box (source code, compiled binaries, static files), and gray-box (partial or full system context). These tests confirmed that frontier AI models achieve high signal-to-noise reasoning against enterprise environments when operating with grounded context.

Scoring Framework

Each organization was scored against the Frontier-AI threat model using a weighted 100-point composite across five defensive pillars: Hide from Attack (25 points, 25%), Limit Blast Radius (25 points, 25%), Govern AI Use (20 points, 20%), Protect the Data (20 points, 20%), and Deception and Decoys (10 points, 10%). Scoring assessed both control ownership and operational activation status.

Data Sources

Assessments drew from multiple categories of input including external attack surface management (EASM), AI-driven vulnerability prioritization, AI-specific surface discovery, private-access configuration reviews, and where possible, inline telemetry covering inspection coverage, policy enforcement, and authentication posture.

Assessment Model

Building on comprehensive EASM visibility, the Opus engine validated and prioritized findings to surface confirmed exploitable issues requiring action, refining the candidate set by 84% on average, and by 92% in the most complex case. It also re-scored CVEs by remediation ROI, cross-referenced findings with CISA KEV, and identified multi-step exploit paths across the validated attack surface.

What This Assessment Covers (and Does Not)

The assessment evaluates external attack surface exposure and security controls’ configuration state. It does not assess SOC detection efficacy at the SIEM rule level, incident response readiness, identity provider security posture, or endpoint detection coverage. The score reflects the likelihood that an AI-powered adversary could penetrate, traverse, and exfiltrate without encountering an effective control. It does not represent a total security posture score.

Critically, frontier-AI-class attacks can compress the kill chain from reconnaissance to exfiltration into minutes, below the response threshold of even mature SOC’s and fast enough to exploit configuration gaps before human intervention.



Key Findings

Finding 1

Enterprises can maximize their existing AI-defense capabilities with few simple steps

Organizations can significantly improve their defenses with practical cyber hygiene.

The activation opportunity is broad. In 64% of organizations, licensed deception platforms are positioned for deployment to add high-fidelity detection depth. 15% of organizations have Intrusion Prevention Systems ready to transition from monitoring to full enforcement mode. 18% have authentication capabilities that can be extended to additional sites. Through configuration optimization alone, organizations can reach target scores between 73 and 98 out of 100. From the cohort’s average starting score of 37, this represents massive individual point improvements.

Why this matters against AI-class threats: Security tools operating below full enforcement provide reduced resistance to an adversary at machine speed.

Finding 2

AI governance is at year zero across every assessed enterprise, regardless of organization size or control breadth

This finding is universal. 100% of assessed organizations lack identity-based authentication for their AI, and 91% have publicly exposed AI endpoints. The average AI governance pillar score was 2.1 out of 20 (10.5% of maximum). The highest score in the cohort was 3 out of 20.

The exposure is substantial: public AI endpoints were common across the cohort, including internet-exposed MCP endpoints and unauthenticated chatbots or copilots. The most serious cases observed included critical-severity GenAI gateways exposed to the public internet.

This gap exists because AI surfaces are deployed by development teams, business units, and third-party vendors at a pace that outstrips security program adaptation. AI governance is not a future initiative. It is a current, exploitable gap in every organization assessed and the fastest-growing category of unmanaged attack surface.



Finding 3

Limited encrypted traffic inspection reduces the effectiveness of deployed security controls across the majority of traffic

Encrypted traffic inspection has long been a security operations priority. In the context of AI-powered threats, full inspection coverage becomes critical. Frontier AI models conducting exfiltration will almost certainly use encrypted channels, particularly AI API calls, cloud storage uploads, and model-to-model communications that blend with legitimate traffic patterns. An adversary operating within uninspected traffic is not visible to inline controls. This elevates encrypted traffic inspection from an operational best practice to a foundational requirement against AI-class threats.

The average enterprise inspects approximately 36% of its encrypted traffic. No organization meets the 70% qualified inspection threshold defined in this assessment as the proportion of encrypted traffic that passes through active inline security controls (DLP, IPS, sandboxing, AI egress) sufficient for those controls to achieve optimal security coverage. In the worst case, 99.6% of all encrypted traffic passes through security infrastructure completely uninspected.

When encrypted traffic remains uninspected, inline security controls like DLP, IPS, sandboxing, and AI egress governance are rendered ineffective, leaving organizations blind to critical threats like data exfiltration and malware delivery. Given that over 86% of cyberthreats and all AI-model interactions use encrypted channels, comprehensive traffic inspection is a non-negotiable prerequisite for defending against AI-class attacks.

Root causes are operational, not technical. The vast majority of organizations use an opt-in inspection model (defaulting to “don’t inspect”). Across the cohort, organizations frequently rely on ‘Bypass Other Policies’ or ‘Do Not Inspect’ SSL rules that disable inspection for entire traffic categories. Among affected organizations, the number of these bypass rules ranged from 2 to 51.



Finding 4

Assessed environments are not yet configured to detect or prevent AI-powered attacks at the speed required

Finding 4 highlights a critical detection gap where enabled security controls fail because their configurations – across deception, IPS, anomaly detection, and session monitoring – are not optimized for AI-class attacks, leaving organizations exposed to threats they are equipped to block but unable to identify. Consequently, a Frontier AI-class attack remains invisible to the SOC, with no triggers for deception, IPS, behavioral anomalies, or session timeouts, meaning the first indication of compromise is likely an external notification rather than internal detection.

- **Deception (the early warning system):** 64% of organizations have deception platforms with coverage not yet extended to critical segments. An attacker moving laterally would not encounter tripwires in the current configuration. Deception provides effective, signature-independent detection against AI-driven attacks by using high-fidelity decoys to trigger immediate, automated containment of lateral movement.
- **Intrusion Prevention (the inline block):** For 15% of organizations, there is a clear opportunity to strengthen defenses by transitioning their Intrusion Prevention Systems into full enforcement mode. Where enabled, “Bypass Other Policies” rules meant most traffic was never evaluated. A known exploit transiting encrypted channels would pass uninspected and unblocked.
- **Application-layer anomaly detection (the behavioral signal):** Multiple organizations had not yet configured inline rules. No anomaly detection on wildcard traversal volume. A single identity accessing thousands of applications in sequence would very likely produce no alert.
- **Session and identity monitoring (the containment trigger):** Authentication timeouts configured without expiration in multiple organizations mean that compromised sessions could remain valid indefinitely. In one organization, none of 18 access rules carried device posture checks. A stolen credential could operate for weeks without triggering re-authentication.
- **What a Frontier AI-class attack very likely looks like to the SOC:** The SOC sees nothing. No deception trigger. No IPS alert. No behavioral anomaly. No session timeout. The first indication of compromise would very likely be external notification rather than internal detection.





Finding 5

Broad segmentation has converted private access platforms into lateral movement freeways

In our analysis, 64% of organizations using ZTNA had wildcard-heavy configurations that may expose more applications than intended. In one case, just 195 wildcard domains made more than 108K applications reachable potentially allowing a single identity to access a very broad portion of the environment. While often introduced during deployment or migration to simplify onboarding, these configurations can persist as environments grow, expanding access beyond original intent. Without regular review and refinement, wildcard-based policies may undermine least-privilege goals and increase the impact of a compromised identity, such as through phishing or session theft.

Multiple organizations compound this with authentication timeouts set to “never.” A compromised session remains valid indefinitely. Even the highest-scoring organization on Limit Blast Radius (14.4 out of 25) still maintained wildcard configurations. Organizations scoring lower had effectively no segmentation.

This pattern means the private access environment does not yet reflect the least-privilege segmentation that zero trust architectures are designed to enforce. Against a Frontier-AI-class adversary capable of identity compromise at scale, wildcard-dominant configurations represent the highest-leverage lateral movement path in assessed environments.

Finding 6

More than one-third of enterprises maintain known exploited vulnerabilities on internet-facing infrastructure

These exposures often persist not because they are unknown, but because remediation must compete with operational realities such as maintenance windows, change-control requirements, and infrastructure dependencies. More than half of the cohort (53%) received remediation recommendations for CISA Known Exploited Vulnerabilities (KEVs) on internet-facing assets. During active probing, more than one-third of enterprises maintained these actively exploitable KEVs on their perimeter.

The result is a growing mismatch between how quickly defenders can remediate and how quickly attackers can identify and exploit weaknesses. As adversaries increasingly use AI to accelerate vulnerability discovery and exploitation, that window of exposure is likely to compress further.

Three CVEs appeared across the majority of the cohort. All are on the CISA Known Exploited Vulnerabilities catalog. All were confirmed via active probing.

- **CVE-2024-38475 (Apache mod_rewrite):** EPSS score 0.9988 to 0.9997
- **CVE-2021-40438 (Apache mod_proxy SSRF):** EPSS score 0.9999 to 1.0000
- **CVE-2023-44487 (HTTP/2 Rapid Reset):** EPSS score 0.999

Their persistence, even in sophisticated organizations, reflects an operational reality: web-tier infrastructure, particularly Apache and F5/NGINX, often falls outside standard patch automation and requires manual coordination across change windows.



While current exposure reflects known vulnerabilities, organizations carrying known KEVs on their perimeter also face the risk that exploitation of still-undisclosed vulnerabilities in the same systems could emerge before existing remediation is complete.

Finding 5 further amplifies this risk, as each unpatched internet-facing asset represents a pre-authentication entry point and, in environments where a single identity can reach hundreds of thousands of applications, one exploited CVE can extend access far beyond the affected server.

The priority actions are to maintain continuous asset visibility through external attack surface management, reduce discoverable attack surface through brokered access, shorten detect-triage-patch cycles for business-justified exposed assets through exploitability-based prioritization, and use autonomous virtual patching during remediation windows.



The New Attack Chain: How Frontier AI Autonomously Chains Small Gaps Into a Single Exploit Sequence

These findings are not isolated. They are reinforcing links in a broader exploit chain. Individually, each is an addressable gap; together, they create materially greater exposure than any one finding suggests.

A frontier AI model with grounded context can autonomously identify and sequence conditions such as an unpatched CVE, IPS not yet in enforcement, wildcard configuration, ungoverned AI endpoint, and uninspected egress path into a continuous exploit path at machine speed, without human involvement. Traditional scanners and penetration tests do not surface this compound path at comparable speed. The following sequence was validated across assessed environments:

Step 1: Enumerate an exposed VPN appliance (one of 500+) via internet scan. Identify CVE-2024-38475 (Apache mod_rewrite, EPSS 0.9997) on the web tier.

Step 2: Exploit the pre-authentication vulnerability. Achieve initial access without credentials. IPS is disabled. No alert is generated.

Step 3: Harvest credentials from compromised VPN infrastructure. Obtain a single valid enterprise identity.

Step 4: Authenticate to the private access platform. Wildcard configuration grants access to 17,000+ applications. Timeout is set to “never.” No device posture check.

Step 5: Traverse applications at machine speed. No deception decoys in path. The SOC sees no alert.

Step 6: Reach an ungoverned AI inference endpoint (one of 300+) accessible without additional authentication. Query proprietary data through the model’s context window.

Step 7: Exfiltrate data via encrypted channel. Two thirds of traffic is uninspected. DLP cannot see the egress. No alert is generated.


Total time from initial scan to exfiltration:

During internal adversarial testing, enterprise AI systems reached a median time to first critical failure of just 16 minutes.

In Zscaler’s frontier model testing program

(Frontier-AI-class and GPT 5.5 across three modes: zero internal knowledge, source code access, and partial/full system context), this chain of reasoning capability was validated. The grounded context harness enables frontier models to identify multi-step exploit paths that neither traditional scanners nor human analysts would surface at equivalent speed.

Every component was observed in organizations’ environments during H1 2026. The chain represents the adversary model against which the Frontier-AI-Readiness Score is calibrated.





Case Evidence: **Three Organizations in Progress**

The following anonymized case studies represent organizations at different stages of their remediation journey. They illustrate starting-state conditions and the operational methodology used to close gaps identified during assessment.


CASE 1

Global Financial Services Organization

Starting state: 10,000+ internet-facing assets. 4,000+ validated CVEs (900+ unique, 165 Critical). Three of the cohort's most prevalent CVEs (CVE-2024-38475, CVE-2021-40438, CVE-2023-44487) confirmed on exposed infrastructure.

Assessment findings: The organization's large internet-facing estate contained critical unpatched vulnerabilities accumulated over successive infrastructure generations. SSL inspection had not yet been extended to cover the majority of traffic. Deception capabilities were entitled but had not reached deployment. AI surfaces remained outside the governance perimeter.

Remediation initiated: Prioritized patching of 165 Critical CVEs. SSL inspection conversion from opt-in to opt-out. Deception activation on crown-jewel financial segments. AI surface gating initiated.





CASE 2

Global Food Producer

Starting state: 13 VPN appliances exposed. 40 public AI surfaces outside the governance perimeter (the largest AI exposure density relative to size in the cohort). 2 validated CVEs on exposed infrastructure.

Assessment findings: The organization’s rapid digital transformation had outpaced the security program’s ability to incorporate new surfaces. AI endpoints had been provisioned by business units as part of operational initiatives and had not yet been routed through the security architecture. Segmentation policy had not been established for the private access environment. VPN infrastructure remained internet-facing pending migration.

Remediation initiated: VPN shielding behind private access (immediate). AI surface gating for 40 endpoints. Segmentation policy creation from zero. Deception deployment on production segments.

CASE 3

Global Luxury Goods Conglomerate

Starting state: 487 public-facing names in scope. 922 TB monthly uninspected encrypted traffic. 10,000 defined applications.

Assessment findings: The organization’s inspection architecture had been deployed with a conservative opt-in model, leaving the majority of encrypted traffic unexamined. With 10,000 applications and near-zero segmentation, the environment had scaled without corresponding access governance.

Remediation initiated: SSL inspection conversion as highest priority. Wildcard replacement across 10,000 applications. Deception activation. AI governance baseline establishment.

Implementation Realism: What the 90-Day Sprint Actually Requires

The path to resilience is technically achievable within 90 days because most often the controls are already owned and dependency chains are well-understood. However, since it is not effortless, it needs to be an executive mandate. Closing the frontier-AI readiness gap is an enterprise initiative that benefits from board-level visibility, CISO to CEO alignment, and dedicated execution capacity. It is an opportunity for shared alignment in an evolving threat landscape. Security leaders should:

1. Brief their boards on their current readiness posture, framing it within SEC, DORA, and NIS2 regulatory context.
2. Secure executive sponsorship for a 90-day sprint with named ownership and regular C-suite reporting.
3. Allocate small, dedicated teams to perform implementation with the expectation that friction encountered due to expanding inspection coverage, access refinements during wildcard replacement, and the like, are an anticipated part of the activation process.

SSL inspection requires cross-functional alignment. Organizations should coordinate work with IT and security teams to fully map current inspection coverage, identify which traffic streams demand inspection, and determine where inspection is operationally infeasible (e.g., certificate-pinned applications or legacy integrations). Where inspection cannot be applied, mitigating controls such as browser isolation, endpoint DLP, CASB-based API inspection, and DNS-layer filtering can provide compensating visibility and threat prevention without requiring inline decryption.

Wildcard replacement requires application owner coordination. Replacing 200+ wildcards with named FQDNs is not a security team exercise alone. Each wildcard resolves to applications owned by business units, development teams, and vendors who must validate access continuity. The security team defines target architecture, but execution requires cross-functional coordination.

The timeline assumes dedicated execution capacity. Organizations that assigned a named sprint team (2 to 4 engineers) progressed on schedule. Those pursuing remediation within existing workloads benefited from additional support to maintain momentum. Their success hinged on executive sponsorship, including regular C-level visibility, a dedicated team with protected capacity, and weekly milestone reporting against the target score.

While any significant security update may cause temporary operational changes, these steps are essential for building a more resilient long-term defense. Our assessment highlights that proactively addressing gaps, such as the initial score of 37, is a critical step in managing modern digital threats. Balancing these short-term adjustments with the need for robust protection is a key part of strengthening the organization's overall security posture.



Recommendations

1

Enable enforcement mode on security controls

Every assessed organization had high-impact controls (IPS, authentication, malicious IP blocking, domain fronting protection) that can be brought to full enforcement through configuration changes. Security operations teams should audit all owned capabilities for activation status, leveraging vendor configuration audits and activation support as appropriate, and enable them in enforcement mode on a 14-day sprint.

We recognize the concern every SOC team has: moving IPS from monitor to block can create false positives and disrupt traffic. The right approach is a phased rollout: start in non-critical segments, tune and validate, then expand to production with clear rollback criteria.

2

Gate all AI surfaces with identity-aware access immediately

Every AI endpoint reachable from the public internet must be placed behind enterprise identity verification. This is recommended for 100% of assessed organizations. Until identity-aware access is in path, AI infrastructure remains exploitable by any internet-connected adversary.

3

Raise SSL/TLS inspection to a minimum of 70% qualified coverage within 90 days

Reverse the model from opt-in to opt-out: inspect all traffic by default, then create narrow exclusions. Eliminate “Bypass Other Policies” rules. Reduce Do-Not-Inspect exceptions to the minimum required. Plan for 2 to 4 weeks of application breakage. Assign helpdesk surge capacity. Without this coverage, DLP, IPS, sandboxing, and AI egress controls cannot function on the majority of traffic.

SSL inspection should be risk-based, not universal. Start with the most critical applications and highest-risk traffic, then expand selectively. Because some applications and services require exceptions, such as certificate pinning or due to regulatory constraints, organizations should define exception paths and rapid response procedures in advance. Where inspection is not feasible, reduce visibility gaps with compensating controls.

4

Replace wildcard domains with named FQDNs on critical segments

Audit all private access configurations for wildcard entries. Prioritize replacement on crown-jewel segments: financial systems, IP repositories, identity infrastructure. Implement authentication timeouts, device posture validation on every access rule, and application-level segmentation limiting each identity to required applications only. Plan for cross-functional coordination with application owners.

5

Deploy deception on critical network segments

Organizations can strengthen their defenses by placing decoys near crown-jewel applications, creating an independent detection layer with high-fidelity alerts that help SOC teams respond faster. When a decoy is triggered, Deception automatically revokes the compromised user’s access to private applications through ZPA across every signed-in device, without requiring analyst approval. This cuts off attackers before they can pivot, stopping lateral movement at the speed of the attack, not the speed of the SOC queue.

6

Move internet-exposed VPN appliances behind zero trust brokered access

Every discoverable VPN appliance is a pre-authentication attack surface. Front all VPN infrastructure with brokered access within 30 days. Decommission exposed appliances on a 90-day timeline. This eliminates the entry point a frontier AI model would probe first.

7

Patch CISA-KEV vulnerabilities on internet-facing assets within 72 hours of assessment

CVE-2024-38475, CVE-2021-40438, and CVE-2023-44487 have EPSS scores at or near 1.0 with confirmed active exploitation. Patch or virtually patch immediately. Where maintenance windows are required, deploy autonomous virtual patching to block exploit attempts in the interim.





Regulatory_

Exposure Note

Organizations scoring below 40 may have control weaknesses that increase both the likelihood and potential impact of a significant cyber incident. We encourage organizations to review and act on the recommendations in this report to help reduce risk exposure and strengthen resilience. In some cases, an incident associated with these weaknesses could also create disclosure considerations under SEC Item 1.05 (Form 8-K).

For organizations subject to DORA or NIS2, missing foundational controls may also point to opportunities to strengthen ICT risk management and overall operational resilience.

Next_

Step

The Frontier-AI-Readiness Assessment provides each organization a quantified starting score, a prioritized remediation sequence, and a 90-day target achievable primarily through expanding or fine tuning existing capabilities. The data in this report establishes the industry baseline. The question for every security leader is where their organization falls relative to it.

You've seen the findings. Now see your exposure.

Talk with a Zscaler expert to identify where AI risk may exist in your environment and what to address first.

Know the risk. Act Fast. Stay Secure.

Learn more at zscaler.com/campaign/zscaler-ai-security-demo

About_

ThreatLabz

ThreatLabz is the security research arm of Zscaler. This world-class team is responsible for hunting new threats and ensuring that the thousands of organizations using the global Zscaler platform are always protected. In addition to malware research and behavioral analysis, team members are involved in the research and development of new prototype modules for advanced threat protection on the Zscaler platform, and regularly conduct internal security audits to ensure that Zscaler products and infrastructure meet security compliance standards. ThreatLabz regularly publishes in-depth analyses of new and emerging threats on its portal, research.zscaler.com.

Follow us: X [@ThreatLabz](https://twitter.com/ThreatLabz) | ThreatLabz [security research blog](https://research.zscaler.com)



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 150 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at [zscaler.com](https://www.zscaler.com) or follow us on Twitter [@zscaler](https://twitter.com/zscaler).

© 2026 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at [zscaler.com/legal/trademarks](https://www.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](https://www.zscaler.com)