




How Zscaler Helps You Build a Mythos-ready Security Program



Anthropic triggered a seismic reaction across cybersecurity when it published the capabilities of its Mythos model, citing its ability to discover and exploit vulnerabilities at machine speed. Anthropic and OpenAI have taken measures to give security teams a headstart through programs like Project Glasswing and the Trust Access for Cyber program, before bad actors can access the models. But — as the Cloud Security Alliance (CSA) detailed in its [expedited strategy briefing](#) — frontier AI models will ultimately provide asymmetric benefits to attackers. In anticipation of the flood of vulnerability disclosures and exponentially shorter times-to-exploit, CSA identified 11 Priority Actions security teams can take to improve their defenses. Eighty CISOs contributed to the CSA's strategy briefing, and another 250 reviewed the recommendations.

This document reviews the 11 recommended actions and how Zscaler products help security teams take fast action in each area.





Table of Contents

Priority Action #1: Point Agents at Your Code and Pipelines 04

Priority Action #2: Require AI Agent Adoption 05

Priority Action #3: Defend Your Agents 07

Priority Action #4: Establish Innovation, Acceleration Governance 08

Priority Action #5: Prepare for Continuous Patching 09

Priority Action #6: Update Risk Models and Reporting 10

Priority Action #7: Inventory and Reduce Attack Surface 12

Priority Action #8: Harden Your Environment 14

Priority Action #9: Build a Deception Capability 16

Priority Action #10: Build an Automated Response Capability 17

Priority Action #11: Standup VulnOps 19



PRIORITY ACTION #1

Point Agents at Your Code and Pipelines

Priority Action overview:

CSA's first recommendation is to turn AI inward and deploy coding agents to review in-house developed code before attackers use theirs to find what you missed. Adopt agent-assisted review on every commit, shift agents left into developer IDEs, and make LLM-driven review a merge gate for all code, whether human or AI-generated.

The hard part is deploying these agents safely. Coding agents like Cursor and Claude Code introduce new traffic (prompts, tool calls, MCP connections) that legacy SWG and CASB controls weren't built to inspect, and they create new exfiltration paths where source code, secrets, and regulated data can leak into external prompts in seconds. Pointing agents at your code improves security only if the agents, their traffic, and the AI components they touch are governed.



HOW ZSCALER HELPS

AI Security

Zscaler AI Protect secures the AI-assisted development environment CSA's first priority depends on: the coding agents, the traffic they generate, and the AI supply chain feeding modern pipelines.

AI Asset Management gives security teams endpoint visibility into AI apps, plugins, and coding agents like Cursor and Claude Code, and it continuously scans source code repositories (GitHub, GitLab) to map

Risk Level:

Critical

CSA Description:

Turn agents and LLM capabilities inward on your own code and dependencies. Start immediately by asking an agent for a security review of any code, then build toward a full audit within your CI/CD pipeline, and shift left by adding capabilities directly into developers' coding agents. All code (human or AI-generated) should pass LLM-driven security review before merge.

Zscaler Products:

AI Security, ZIA, ZPA

AI workflows, MCP servers, and supply-chain risks from AI packages and open-source components.

Secure Access to AI governs how those agents reach public LLMs and private apps — with inline inspection, allow/block/isolate/caution policies, and DLP that recognizes source code and secrets as sensitive.

AI Guard adds runtime guardrails on prompts and responses, blocking prompt injection and jailbreaks and preventing credentials, secrets, PII, and source code from leaking in either direction.

Together, these controls let teams safely expand coding-agent adoption knowing the agents, their traffic, and their outputs are continuously governed.

Zscaler Internet Access (ZIA) and Zscaler Private Access (ZPA)

Zscaler Internet Access can protect MCP traffic by treating it as outbound Internet / API traffic and applying inline threat and data protection controls that are used for any SaaS / API access, particularly when MCP is carried over HTTPS or WebSockets.

Zscaler Private Access (ZPA) can secure agent API calls to private apps on a per-application basis.



PRIORITY ACTION #2

Require AI Agent Adoption

Priority Action overview:

CSA is intentional in using the word “require.” Optional AI adoption programs fail to overcome cultural friction, and “human-speed” security operations are a losing strategy against attackers running parallel kill chains in minutes. Every downstream priority — continuous patching, automated response, VulnOps — assumes analysts operating at agent speed.

Required adoption also requires guardrails. The same agents that accelerate triage, code review, and remediation will handle sensitive data and call privileged APIs. Security leaders need mandated adoption to mean sanctioned models, enforced policy on agent traffic, DLP on what agents can say and see, and audit trails that satisfy regulators and the board.



HOW ZSCALER HELPS

AI Security

AI Protect makes mandating agent adoption operationally defensible. **AI Asset Management** maintains a unified inventory of AI apps, models, agents, and MCP servers across users, infrastructure, and code — including shadow AI — and supports model approval workflows so adoption happens on sanctioned LLMs. Endpoint telemetry provides evidence of AI deployment — including coding agents — to provide an audit trail for sanctioned agents. **Secure Access to AI**, enforced through the Zero Trust Exchange, applies allow/block/isolate/caution policy per user, group, and app; restricts access to corporate AI tenants while blocking personal

Risk Level:

Critical

CSA Description:

Formalize AI agent usage (mostly in the form of “coding agents”) as part of all security functions, with mandatory security controls and oversight in place. While defensive AI technology has not yet caught up, these agents empower staff to be effective in the new threat landscape, allowing acceleration beyond “human speed.” Optional adoption programs have not been shown to overcome cultural barriers, while adoption is a limiting factor in achieving the rest of the actions in this table.

Zscaler Products:

AI Security, Agentic SecOps


accounts; and uses inline DLP to prevent credentials, secrets, PII, and source code from leaving the enterprise in prompts. By enforcing “responsible use of AI,” security leaders create a culture where the organization is comfortable with AI agent adoption.

AI Guard provides the runtime guardrails that make broad adoption safe: intent-based detection for prompt injection, jailbreak, toxicity, and off-topic behavior, with stateful multi-turn attack detection and enforcement actions (block, isolate, rate limit) applied consistently across public GenAI and agent workflows. A full audit trail and Conversational Viewer produce the governance-ready evidence CSA is calling for, so CISOs can require agent adoption across the SOC, IT, and engineering with confidence that moving fast to adopt agents isn’t creating new risk faster than it’s retiring old risk.

Zscaler provides the risk assessment, guardrails, and oversight that allow security leaders to require AI agent adoption.



FIGURE 1: Deploy AI Agents Securely with Runtime Guardrails



Agentic SecOps

To meet CSA’s recommendation to require AI agent adoption, **Agentic SecOps** helps SOC teams standardize agentic workflows for triage, investigation, and response so they can operate at the speed Mythos-era threats demand. It unifies alerts — from Zscaler and third-party tools — into prioritized threat stories; enriches them with identity, asset criticality, and exposure context; and uses AI-generated summaries and attack-path mapping to help analysts quickly understand which threats matter and what they should do next.

Agentic SecOps helps secure agent adoption through oversight and repeatability. SOC teams can apply business-aligned grouping and prioritization logic, require human approval for sensitive steps, and capture evidence and actions in a consistent workflow while still enabling fast, right-sized containment actions when seconds count.

Managed detection and response (MDR) extends mandated AI agent adoption into the SOC for organizations that need proven, production-hardened agentic workflows without the build time. AI agents perform discrete investigative tasks as part of MDR: gathering telemetry, enriching alerts with IP reputation, geolocation, and behavioral context, and surfacing tailored remediation recommendations built on years of historical threat data. Investigations that would take a human analyst 30+ minutes are finished in under three minutes, with MDR detection engineers continuously validating and optimizing for true positives. For organizations looking to formalize AI agent adoption, MDR provides a proven SecOps model that already runs at scale.





PRIORITY ACTION #3

Defend Your Agents

Priority Action overview:

As agents move into production, they become high-value targets. Existing controls weren't built to reason about prompts, tool calls, retrieval pipelines, or escalation logic, and they miss the failure modes that matter most: prompt injection, jailbreaks, context poisoning, overprivileged tool use, and compromise of the models, packages, and MCP servers agents depend on. CSA is explicit that the agent harness — prompts, tool definitions, retrieval pipelines, escalation logic — is where the most consequential failures occur.

CSA is equally explicit that organizations should not wait for external governance frameworks. Defining scope boundaries, blast-radius limits, and human override mechanisms — and running continuous adversarial testing against the harness — is a design problem security teams own today.



HOW ZSCALER HELPS

AI Security

AI Protect defends agents across build, deploy, and run. **AI Asset Management**, which provides security posture management for AI, maintains an AI BOM of every LLM, agent, workflow, MCP server, and supporting data resource, with model lineage, supply-chain risk scoring, MCP server discovery and hardening, and cross-layer risk mapped to frameworks including NIST AI RMF, EU AI Act, ISO 42001, OWASP LLM Top 10, and MITRE ATLAS.

Risk Level:

Critical

CSA Description:

Without agents, most tasks on this list will be untenable, but they must be defended. Agents are not covered by existing controls and introduce both cyber defense and agentic supply chain risks. The agent harness — prompts, tool definitions, retrieval pipelines, and escalation logic — is where the most consequential failures occur; audit it with the same rigor as the agent's permissions. Before deploying agents in or adjacent to production environments, define scope boundaries, blast radius limits, escalation logic, and human override mechanisms. Do not wait for industry governance frameworks. Define your own now.

Zscaler Products:

AI Security

AI Red Teaming audits the agent harness with the rigor CSA demands: automated adversarial probes across APIs, major LLM providers, conversational platforms, and AI development tools, with CI/CD-triggered assessments, remediation guidance, system-prompt hardening, and LLM benchmarking for model selection. Findings can feed directly into runtime guardrail policy.

AI Guard delivers runtime defense: intent-based detectors, prompt injection and jailbreak blocking, data-leakage prevention for credentials, secrets, PII, and source code, and stateful multi-turn attack detection with block/isolate/rate-limit enforcement.

Together, these capabilities give security teams the four elements CSA requires to defend agents: scope, harness testing, runtime policy, and containment.



PRIORITY ACTION #4

Establish Innovation, Acceleration Governance

Priority Action overview:

CSA’s fourth priority is the organizational scaffolding that makes everything else possible. Security, Legal, and Engineering have to work together to evaluate new offensive threats and onboard defensive technologies at attacker speed. The failure mode is familiar: review processes designed for quarterly software purchases become the bottleneck that hands attackers their asymmetric advantage.

Effective governance needs a continuously updated picture of AI use inside the enterprise — what is being used, by whom, on what data — and a defensible policy and testing framework so onboarding new AI capabilities doesn’t require re-litigating first principles each time.

 HOW ZSCALER HELPS

AI Security

AI Protect gives the cross-functional governance team a single source of truth and a repeatable framework for accelerating approvals. **AI Asset Management** continuously inventories AI apps, models, agents, MCP servers, and supply-chain components across endpoint, internet, cloud, SaaS, and code repositories — turning “what AI do we have?” from a quarterly survey into a live dashboard. Compliance posture management benchmarks that inventory against 35+ prebuilt frameworks including

Risk Level:
Critical

CSA Description:
Cross-functional mechanism (Security, Legal, Engineering) to evaluate new offensive threats and accelerate onboarding of defensive technologies. Without this in place, every other action in this table runs into approval friction that slows deployment to the attacker’s advantage.

Zscaler Products:
AI Security

NIST AI RMF, EU AI Act, ISO 42001, OWASP LLM Top 10, and MITRE ATLAS, giving Legal and GRC defensible mappings without custom work.

AI Red Teaming gives governance a standardized gate: any new AI app, agent, or workflow can be probed against adversarial tests, with results, remediation guidance, and compliance mapping produced automatically and triggerable from CI/CD. **Secure Access to AI** and **AI Guard** then operationalize approved policy at runtime — allow, block, isolate, or caution at the Zero Trust Exchange, with **DLP** and runtime guardrails enforcing the guardrails governance agreed to. The result is a governance function that evaluates, approves, and monitors AI at the pace the business and the threat landscape demand.

FIGURE 2: Requirements for Effective Governance

	Source of truth “What’s running right now?”	AI Asset Management
	Compliance benchmarks NIST AI RMF, ISO 42001, OWASP LLM Top 10, etc.	AI Asset Management
	Standardized testing gate Probes any new AI asset against adversarial tests	AI Red Teaming
	Safely deploy Operationalize approved AI assets at runtime	Secure Access to AI



PRIORITY ACTION #5

Prepare for Continuous Patching

Priority Action overview:

One objective of [Anthropic's Project Glasswing](#) is for partners such as Amazon Web Services, Apple, NVIDIA, Zscaler, and the Linux Foundation to discover vulnerabilities using Claude Mythos and release patches. It provides defenders with a headstart against attackers, who may soon gain access to frontier AI models capable of exploiting hundreds of vulnerabilities in commonly used products in seconds.

As Anthropic partners discover code flaws and vulnerabilities, security teams should expect a massive influx of findings and available patches. Efficient patch management has always been a goal of proactive security, but the upcoming volume of critical vulnerabilities and associated patches will apply immense pressure on cross-functional patching workflows.

While Zscaler does not offer a solution for patch deployment, it does provide customers with the prioritization and remediation workflow engines required to facilitate efficient patching between SecOps and remediation owners in IT, DevOps, Cloud Security, and other teams.

Risk Level:

Critical

CSA Description:

With the increase in vulnerability discovery and reporting, and specifically now that Glasswing has made Mythos available to significant software vendors, prepare triage and deployment capacity to handle a potential flood of patches as new critical vulnerabilities are disclosed.

Zscaler Products:

Exposure Management




HOW ZSCALER HELPS

Exposure Management

Automated remediation workflows and AI-powered patch recommendations are two staples of Zscaler [Exposure Management](#). Customers create grouping rules to automatically cluster findings by assignee, business unit, required fix, or any other criteria to automate and simplify triage. The Remediation Copilot answers prompts with tailored remediation plans — including optimal patches and software upgrades — and pins suggested fixes to tickets. Finally, bidirectional workflows with popular ticketing tools like ServiceNow and Jira deliver remediation owners the full context associated with the work request and ensures that any action taken in the ticketing platform, for example, is reflected in Exposure Management for dynamic and complete visibility.

FIGURE 3:

Identify the most impactful fixes with AI





PRIORITY ACTION #6

Update Risk Models and Reporting

Priority Action overview:

Frontier AI models are redefining cyber risk, and every organization must do the same. As the volume of vulnerability disclosures increases and speed-to-exploit decreases, security teams need razor-sharp risk modeling to focus their resources. Static scoring like the Common Vulnerability Scoring System (CVSS) and Exploit Prediction Scoring System (EPSS) lack environmental context about your assets, multiplying risk factors such as open ports or misconfigurations, and mitigating controls blocking attack paths. As a result, security teams waste precious time and resources chasing “false criticals,” reporting on generic findings and patches without a perspective on the reduction of actual business risk.

Legacy performance metrics such as vulnerabilities patched, mean-time-to-remediate (MTTR), and SLA compliance have little to do with the risk reduced or the likelihood of an attack. In fact, strong metrics in those areas — presented in isolation — can result in inadequate funding for a new world of AI-powered attacks that won’t wait for patch windows.

Risk Level:

Critical

CSA Description:

Review and update security risk metrics, reporting, and business risk calculations to reflect AI-accelerated exploit timelines and attack complexity. Pre-AI assumptions about patch windows, exploit scarcity, and incident frequency may no longer hold. Outdated models could potentially even lead to underfunding of controls and inaccurate business reporting. Communicate and collaborate with stakeholders, mapping out and prioritizing potential effects on the business, reporting, and projections.

Zscaler Products:

Exposure Management

To keep pace with the volume of disclosures and shrinking exploit windows, risk models must capture the severity of exposures with environmental context. Static scoring measures generic exploitability. Updated risk models should measure exploitability in your environment, mapped against your controls. Contextual risk allows teams to focus their limited patching cycles and, alternatively, close attack paths with mitigating controls when patching is not an option.

Updated models should quantify whether risk is open or closed as opposed to “patched” or “unpatched.” They should account for business context, real-time threat intelligence, and toxic combinations that indicate critical priority. Most importantly, risk calculations can no longer wait for manual reconciliation of signals from siloed tools — it must be near-real-time.



Exposure Management

Zscaler **Exposure Management** allows customers to understand their riskiest exposures in the context of unique risk factors and mitigating controls. Customers can tune the risk model, adding new factors or adjusting the weighting of existing factors. The scoring engine continuously maps findings to Zero Trust protections and other mitigating controls, reducing false criticals and surfacing findings that are exploitable within a unique environment. By accounting for Zscaler controls, security teams understand “reachability” — the degree to which a vulnerability or security flaw can be exploited. The added context directs limited remediation resources toward the riskiest findings.

The reporting engine includes pre-built and customizable reporting templates to measure organizational risk according to the tailored scoring model. Security teams can schedule regular risk reports for specific business units, asset groups, remediation

owners, or any other relevant logic to the organization. As the risk landscape evolves in a post-Mythos world, so will your risk modeling and reporting.

Risk360

Risk360 combines Zscaler telemetry, security research from ThreatLabz, external attack surface monitoring, and third-party risk signals to generate an executive risk profile. When CSA references “funding of controls,” Risk360 provides Zscaler customers with a cyber risk quantification engine capable of evaluating risk posture against security investments at an executive level. It measures risk across the four stages of an attack using more than 140 risk factors, consolidates signals from various sources, and enhances CXO and board-level understanding of cyber risk management, strategy, governance and compliance, and use of cyber insurance to offset risk.




FIGURE 4:

Zscaler Exposure Management prioritizes risk with complete environmental context.



PRIORITY ACTION #7

Inventory and Reduce Attack Surface

Priority Action overview:

CSA is intentional in its word choice of this two-part action: “inventory and reduce.”

Understanding the attack surface is foundational to risk reduction. You can’t protect what you can’t see, and Zscaler customers have a unique ability to inventory their assets with native discovery tools, the Zscaler Client Connector (ZCC) agent, and third-party integrations with key sources of asset data.

Customers also have access to industry-leading Zero Trust solutions to prevent attackers from discovering services. A zero trust architecture makes invisibility a primary control, reducing the attack surface and fulfilling CSA’s priority action.



HOW ZSCALER HELPS

Zscaler Exposure Management

By leveraging telemetry from the ZCC agent along with out-of-the-box EASM capabilities, Zscaler Exposure Management provides customers with a robust asset inventory sourced natively from their existing Zscaler tools. The **Data Fabric for Security** then harmonizes asset and exposure data from key third-party sources, including vulnerability scanners, EDR tools, CSPM tools, CMDBs, and more. The result is a comprehensive, near-real time inventory that covers endpoints, cloud resources, IT/OT/IoT devices, and any other asset in scope.

Risk Level:

High

CSA Description:

Make use of, update, or create an inventory. Using agents, the process can be significantly accelerated and enable continuous updates. Start with critical internet-facing systems, build toward a full-coverage inventory over 45 days. Generate real SBOMs. Aggressively shut down unneeded or unmaintained functionality, phase out suppliers that no longer comply with your updated vulnerability management requirements, and isolate or airgap at-risk systems. You cannot patch, segment, or defend what you don’t know exists.

Zscaler Products:

Exposure Management, ZDX, ZPA, Zero Trust Branch

This risk-based asset inventory gives Zscaler customers the ability to enforce asset security controls, identify and close coverage gaps (such as assets missing EDR), and ensure all relevant assets are included in their exposure management program.

“When you’re working in a very large organization and you have the number of assets that we’re trying to protect, actually having an inventory is the very first step. The Zscaler Data Fabric for Security gives you a shortcut to get there.”

RYAN WINN

CISO, AdventHealth



Zscaler Digital Experience (ZDX)

ZDX provides comprehensive asset and software inventory data from the endpoint, enabling updates and identification of outdated systems.

ZDX addresses the requirements outlined in the priority action in the following ways:

Inventory and Updates: ZDX utilizes the Zscaler Client Connector (ZCC) agent to collect and deliver a near-real-time inventory of assets, which is a foundational requirement for risk reduction. The ZCC agent accelerates the process of maintaining a continuous inventory, providing an excellent baseline for attack surface coverage that can be augmented with insights from third-party sources.

Software Inventory: ZDX offers a consolidated view that includes Device Inventory and Software Inventory capabilities. As a result, IT teams understand the software portfolio and versions deployed across the organization and on each device. This capability is critical to detecting outdated software before it can be exploited, supporting the need to manage unmaintained functionality.

Software Patch Inventory: ZDX tracks both software and security updates, allowing administrators to monitor patch distribution, patch names, types, and installation dates on individual devices via the Software Patch Inventory feature. This integrated system provides a clear overview of software update status across the entire device fleet.

Device Health and Identification: ZDX provides deep insights into device health metrics such as CPU, memory, and disk I/O. This visibility helps IT teams identify problematic or under-provisioned devices that could be running outdated or unmaintained software and therefore pose a security risk.

Device Remediation: ZDX includes an Device Remediation capability that allows IT to proactively resolve these identified issues.

Remote Scripts at Scale: A key component of this functionality is the ability to enable Remote Scripts

to be run on multiple endpoints simultaneously directly from the ZDX portal.

Zscaler Private Access (ZPA)

Zscaler Private Access (ZPA) directly mitigates the exposure created by internet-facing systems by removing the need to publish internal applications to the public internet. With ZPA, private apps no longer require inbound firewall ports, public IP addresses, or externally reachable DNS that can be scanned and targeted. Instead, ZPA uses an inside-out architecture: lightweight app connectors in your environment establish outbound connections to the Zscaler cloud, and authorized users are connected to specific applications through the service. This approach leaves applications effectively non-discoverable to unauthorized parties — reducing the reachability of critical systems, and in turn, the overall risk of the business.

ZPA also eliminates traditional VPN dependence, another common internet-exposed entry point. Rather than granting broad network access, which increases blast radius and enables lateral movement, ZPA provides identity- and policy-based access on a per-application basis, aligning with Zero Trust principles and limiting attacker pathways.

In a Mythos-response hardening plan, ZPA is a critical control for rapid attack surface reduction: it removes public exposure quickly, reduces reliance on legacy remote access, and helps ensure “unknown” applications are not left exposed while inventories and remediation mature.

Zero Trust Branch

For many organizations, the extended attack surface includes branch offices, factories, clinics, and warehouses. Zscaler **Zero Trust Branch** provides a detailed inventory of IoT/OT devices at branch locations and reduces the attack surface by eliminating exposed VPN ports inherent to traditional site-to-site networking and SD-WAN. Your branches and remote sites become invisible to internal and external malicious actors.



PRIORITY ACTION #8

Harden Your Environment

Priority Action overview:

The basic blocking and tackling of “security hardening” is not only still valid but is perhaps more important than ever, as a means of shrinking the enterprise attack surface from high-velocity attacks. Security teams cannot expect to detect and remediate all Mythos-powered attacks before they cause damage. However, taking measures to reduce vulnerabilities, protect against compromise, and limit the blast radius will improve defenses.

These elements are also all foundational to a zero trust strategy.

CSA suggests three areas in particular to consider:



Inspecting traffic with inline security controls



Granularly segmenting your environment and enforcing least-privilege access policies



Reducing the complexity of your application code



HOW ZSCALER HELPS

Zscaler Internet Access (ZIA)

Priority Action #8 emphasizes hardening controls that raise attacker cost and reduce the impact of what gets through. **Zscaler Internet Access (ZIA)** supports this mission by enforcing egress filtering and full inline inspection of internet and SaaS traffic at

Risk Level:

High

CSA Description:

The basics remain valid and can be prioritized for risks that can't be easily mitigated. Implement egress filtering (it blocked every public log4j exploit). Enforce deep segmentation and zero trust where possible. Lock down your dependency chain. Mandate phishing-resistant MFA for all privileged accounts. Every boundary increases attacker cost.

There are also aspects of this which could be accelerated with AI. For example, software minimization is a high leverage function that reduces the operational overhead of second order functions such as patching. For example, minimizing base operating system images, or replacing third-party libraries with framework primitives as they emerge over time. AI can do this.

Zscaler Products:

ZIA, ZPA, Zero Trust Branch

scale, including encrypted TLS traffic. This capability matters because modern threats frequently hide in outbound and encrypted channels, and consistent inspection is required to stop exploitation and data loss attempts that bypass basic controls.

ZIA applies inline, AI-powered security engines such as advanced threat protection and sandboxing to detect known and unknown threats in real time. Policies can be risk- and context-aware (user, device posture, location, application, data sensitivity), producing different outcomes based on conditions—allowing legitimate activity, blocking malicious traffic, or isolating sessions to contain risk while maintaining productivity. This approach directly aligns to hardening: tighter outbound control, stronger content inspection, and automated enforcement that limits attacker options without relying on brittle perimeter rules.

Zscaler Private Access (ZPA)

Zscaler Private Access (ZPA) operationalizes “deep segmentation and zero trust” by replacing broad network access with application-specific access. ZPA creates 1:1 connections between an authenticated user and an authorized application, creating a “segment of one,” which prevents attackers from moving laterally as they can in flat networks or VPN-based access models. Compromise of a device no longer translates into reachability of the rest of the environment.

ZPA also accelerates hardening by removing inbound exposure: private apps do not require reachability from the internet, reducing scanning and exploitation opportunities. In addition, ZPA AppProtection inspects private application traffic for the same classes of threats you’re trying to prevent at the boundary: malicious payloads, exploit attempts, and common web/API attack patterns.


Zero Trust Branch

Zero Trust Branch proxies all north-south traffic from various sites through the Zscaler Zero Trust Exchange, eliminating the need for a routed overlay. An infected device in one location cannot reach your crown jewel apps in the cloud.

Further, Zero Trust Branch also isolates every device into a “network-of-one” using /32 segmentation. All east-west communication between devices is controlled by group policies based on device attributes, and not just network presence.

Zero Trust Branch inspects and protects all user and IoT/OT traffic going in and out of a location using the rich cyber threat protection and data security capabilities in the Zero Trust Exchange. Egress firewall policies can limit external communication to trusted destinations, with adaptive AI-powered threat protection.

FIGURE 5: Industry’s First AI-Powered App Segmentation with ZPA





PRIORITY ACTION #9

Build a Deception Capability

Priority Action overview:

Detection velocity is one of the highest risks in building a Mythos-ready security program according to the CSA's response. The briefing rates "Inadequate Incident Detection and Response Velocity" as CRITICAL — its highest severity designation — because "alert triage volumes, SIEM correlation speed, and containment authorization latency were designed for human-paced threats." Tools like EDR, NDR, SIEM all assume human-speed attacks, correlating events across multiple signals over minutes, hours, or days before generating a high-confidence incident. Against an agentic attacker running parallel attack paths and completing full kill chains in minutes, traditional detection tools are still correlating while the AI-driven attack has already finished.

The CSA's Priority Action #9 calls for a class of detection that is explicitly attack-tool and vulnerability independent. Signature-based detection fails by definition — the vulnerability didn't exist in your detection database an hour ago. Behavioral detection relies on anomalies and does not fare much better because attackers operate with legitimate credentials and valid tools such as PowerShell and APIs.



HOW ZSCALER HELPS

Deception

Zscaler Deception covers an entire environment with decoy, lures, and breadcrumbs layered across every attack surface an adversary would traverse: network decoys in every segment; fake identities in Active Directory; decoy cloud resources in

Risk Level:

High

CSA Description:

Deception is attack-tool and vulnerability independent, identifying attacks and attackers based on their TTPs. Deploy canaries and honey tokens, layer behavioral monitoring, pre-authorize containment actions, and build response playbooks that execute at machine speed.

Zscaler Products:

Deception

AWS, Azure and GCP; lures on endpoints; and even decoy AI infrastructure include chatbots and LLMs. This coverage makes **deception decisive against agentic attacks** that explore exhaustively and probe every path they can find.

A deception signal carries no ambiguity because no one has a legitimate reason to interact with a decoy. The result is more than just a high-fidelity alert — Zscaler Deception can pre-authorize and configure automated responses that execute the moment a decoy fires, to isolate the compromised host, block the IP, revoke the credential, or take other containment actions. These responses happen at the speed of the attack, not relying on a SOC to triage and investigate first, enabling the machine-speed response the CSA calls for. AI-generated attacks are also computationally expensive, and every interaction with a Deception decoy forces the attacker's model into costly reasoning loops that consume real compute and return nothing. Deception doesn't just neutralize the threat; it can degrade the attacker's ROI with every move they make.

For organizations already running Zscaler Private Access, Zscaler Deception can be deployed in a few hours via a single-click enablement, making the CSA's 90-day window to deploy a deception capability an operational reality.

To learn more, read the [full blog post on the CSA's mandate to deploy deception capability](#)



PRIORITY ACTION #10

Build an Automated Response Capability

Priority Action overview:

CSA urges CISOs to move beyond human-speed detection and ticket-driven remediation to develop machine-speed containment. In a Mythos-ready world — where AI can rapidly identify, chain, and exploit weaknesses — security teams can't rely on analysts manually reviewing alerts, debating severity, and then coordinating response steps across multiple teams. The goal is to pre-build and continuously refine automated response playbooks that can take immediate action — for example, isolating a workload, disabling a risky identity session, revoking tokens, blocking suspicious egress, or triggering an environment “kill switch” — the moment high-confidence signals appear.

Of course, autonomous response requires clean telemetry, well-integrated tools that harmonize threat signals, and carefully tuned playbooks that can operate with the highest possible degree of automation without losing control of the outcomes.



HOW ZSCALER HELPS

Agentic SecOps

Agentic SecOps helps customers build automated responses by turning fragmented alerts into unified, context-rich threat stories and then driving consistent triage and containment workflows. It correlates signals across Zscaler and third-party tools; enriches incidents

Risk Level:

High

CSA Description:

Improve detection engineering and incident response capabilities to be systemic and, to the degree possible, autonomous. Examples: asset and user behavioral analysis, preauthorized containment actions, and response playbooks that execute at machine speed.

Zscaler Products:

Agentic SecOps, Exposure Management

with identity, asset, and exposure context; and applies risk-focused prioritization so automated playbooks focus on the threats most likely to impact the business.

To execute response “at machine speed,” Agentic SecOps translates prioritized threat stories into impact-based containment through guided workflows, inline Zero Trust actions, and orchestrated playbooks. Teams can preauthorize severity-based actions to reduce risk while minimizing disruption; trigger native controls to block or unblock URLs, files, and source IPs; and integrate with SOAR and ITSM for automated workflows and bidirectional ticketing. The result is a systemic, repeatable response model where investigations can be completed in minutes and containment can be applied quickly and proportionally as new attacker behavior emerges.

Managed detection and response (MDR) delivers customizable, no-code automated threat containment. MDR playbooks link triggers to discrete containment and communication actions, facilitating responses such as endpoint isolation, IP and URL banning, account suspension, and identity session revocation. Teams can require manual approval for sensitive actions or enable fully automatic execution, reducing mean time to response by up to 10x. Automated responses can be enforced in not only Zscaler Internet Access but




also third party platforms like CrowdStrike, Okta, and Microsoft. MDR also provides the automated user behavior analysis that the CSA specifically calls out, surfacing anomalous activity and delivering the identity context needed for rapid response.

Exposure Management

The same inline Zero Trust actions and orchestrated playbooks available to your SOC will soon be available in the Zscaler Exposure Management solution.

As frontier AI models make machine-speed exploits commonplace, standard “scan-and-patch” vulnerability management processes no longer reduce risk fast enough. While your teams wait for a patch from a third-party vendor, the window for attackers is wide open — unless you reduce the risk with an automated response like isolating the asset or restricting network/application access. As proactive security expands to autonomous risk reduction, Zscaler provides the playbooks and in-line controls needed for a post-Mythos world.

FIGURE 6: Trigger AI-powered responses with Zscaler Agentic SecOps





PRIORITY ACTION #11


Standup VulnOps

Priority Action overview:

The final Priority Action issued by the CSA includes several key phrases. “Staffed and automated like DevOps” suggests that organizations must have appropriate investments in both people and technology to manage an incoming avalanche of vulnerability findings and exposures.

Most organizations are prepared for “discovery of vulnerabilities” with the scanners they have already deployed. Frontier AI models promise to drive more efficient vulnerability discovery in first-party systems and code bases.

The real challenges in vulnerability management are familiar:



The volume of findings will exceed any prior precedent with AI-powered discovery



The speed of exploitation will increase dramatically, approaching zero, resulting in more exploits occurring on Day Zero or even before CVE disclosure

To address the volume of findings and the speed of exploitation, CSA calls for “autonomous vulnerability research, automated remediation pipelines, and triage discipline.”

Risk Level:

High

CSA Description:

Improve detection engineering and incident response capabilities to be systemic and, to the degree possible, autonomous. Examples: asset and user behavioral analysis, preauthorized containment actions, and response playbooks that execute at machine speed.

Zscaler Products:

Agentic SecOps, Exposure Management



HOW ZSCALER HELPS

Exposure Management

Another way to think about “autonomous vulnerability research” is the ability for technology to decipher which vulnerabilities are riskiest in your environment and how to fix them. Zscaler **Exposure Management** determines the reachability and risk of every finding on every asset by building complete context within the environment. With the context of mitigating controls applied to each finding, security teams can understand what’s truly exploitable in their environment rather than chasing “theoretically exploitable” false criticals.

If prioritization is the first half of VulnOps, automated remediation is the second half. Zscaler Exposure Management breaks the outdated manual cycle of scanning, triaging, patching, and confirming. Automated ticket grouping and triage ensure critical findings are assigned and addressed. Remediation copilot suggests fixes that drive maximum risk reduction. Integrated workflows connect seamlessly with existing remediation systems such as custom ServiceNow deployments. In addition, its tight integration with the Zero Trust policy engine provides the inline controls and actions required to reduce risk at machine speed, satisfying the CSA’s call for automated response capabilities.



Summary: Complete Mapping of Zscaler Products to CSA Priority Actions for Mythos–Readiness

The verdict is clear: Frontier AI models can find and exploit vulnerabilities faster than humans. Mythos Preview proved capable of unprecedented discovery and exploitation, and it’s a matter of time before bad actors gain widespread access to similar technology. Even as trusted partners gain early access to frontier AI models to find and fix vulnerabilities and code flaws, the volume of work for security teams is quickly snowballing beyond the capacity of reactive response executed entirely by humans.

Zscaler is advising customers to follow the guidance laid out by CSA in its executive briefing, prioritizing the 11 actions summarized in this document. This table recaps how your Zscaler investments map to each CSA Priority Action:

CSA Priority Action	Zscaler Product Coverage	Level of Coverage
Point Agents at Your Code and Pipelines	AI Security, ZIA, ZPA	🕒
Require AI Agent Adoption	AI Security, Agentic SecOps, Agentic Framework	🕒
Defend Your Agents	AI Security	🟢
Establish Innovation, Acceleration Governance	AI Security	🟢
Prepare for Continuous Patching	Exposure Management	🕒
Update Risk Models and Reporting	Exposure Management, Risk 360	🟢
Inventory and Reduce Attack Surface	Exposure Management, ZDX, ZPA, Zero Trust Browser	🟢
Harden Your Environment	ZIA, ZPA	🟢
Build a Deception Capability	Deception	🟢
Build an Automated Response Capability	Agentic SecOps, Exposure Management	🟢
Standup VulnOps	Exposure Management	🟢

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 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 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.



Act Fast.
Stay Secure.™