THE PRIORITY INVERSION

Why the SANS Mythos Report Has the Order Wrong

An Executive Perspective

AVIATRIX®

Executive Summary

In April 2026, sixty of the most respected cybersecurity experts in the world published the SANS Mythos Report, eleven priority actions for an era of AI-accelerated threats. Among those actions, “Prepare for Continuous Patching” was rated CRITICAL. “Harden Your Environment,” the construction of architectural constraints that limit damage independent of whether a vulnerability has been patched or a breach has been detected, was rated HIGH.

This paper presents evidence that the priority ordering is inverted.
A companion paper, The Vulnerability Deficit: Why Remediation Cannot Outrun Discovery, demonstrates mathematically that remediation has a structural ceiling while discovery compounds exponentially. A 6.5x increase in remediation effort across more than ten thousand organizations produced worse outcomes, not better ones. The percentage of critical vulnerabilities unresolved at seven days rose from 56% to 63% even as organizations closed 6.5 times more tickets.

The implication is direct. A defense strategy that prioritizes patching speed over containment architecture is optimizing the wrong variable. This paper makes the case using eight structural axioms and evidence from the very threat landscape the Mythos report describes. The conclusion is not that the SANS report is wrong. It is that one of its priority actions must be elevated from HIGH to CRITICAL, and that the evidence for doing so is not a matter of opinion. It is a matter of mathematics.

The Report That Got It Almost Right

The SANS Mythos Report is the most comprehensive analysis of AI-accelerated cyber threats published to date. Its sixty-plus authors and eighty-plus CISO reviewers produced a document that correctly identifies the structural shift: AI has fundamentally changed the attacker-defender dynamic, and the cybersecurity industry must respond at an architectural level, not merely an operational one.

The report’s eleven Priority Actions are, individually, sound. The question this paper raises is not whether these actions are correct but whether they are correctly ordered.

Priority ordering matters because organizations allocate resources in sequence. A CISO reading the Mythos report will see Continuous Patching rated CRITICAL and Hardening rated HIGH. The rational response is to fund patching programs first and hardening programs second. That sequencing decision, replicated across thousands of enterprises, will determine how much damage the next wave of AI-accelerated attacks actually causes.

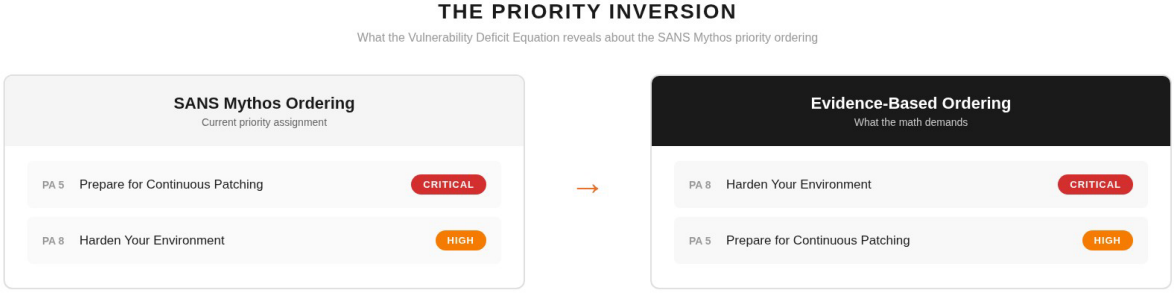

Figure 1: The Priority Inversion

The Toxic Combination Compounds the Math

Paper 1 in this series introduced the Toxic Combination. Three structural forces are converging. Attackers are industrializing. AI is putting frontier offensive capability into millions of hands at consumer cost. Cloud is insecure by default by design. Each leg accelerates the others.

The priority ordering question must be answered against this landscape, not against the threat model of five years ago. A patching strategy that could keep pace with the discovery curve of 2020 cannot keep pace with the discovery curve of 2026, because AI has made offensive capability emergent, permanent, and democratized. A patching strategy that could keep pace with a vulnerability surface bounded by what skilled human operators could identify cannot keep pace with a surface tested by orders of magnitude more actors at consumer compute cost. The math has gotten worse, faster, than the priority ordering reflects.

The Vulnerability Deficit

The full mathematical treatment appears in the companion paper, The Vulnerability Deficit: Why Remediation Cannot Outrun Discovery. The core finding is summarized here because it is the foundation of the priority inversion argument.

The stock of exploitable, unpatched vulnerabilities in any environment is governed by a simple relationship: new vulnerabilities enter the system through discovery and new code, while vulnerabilities leave the system through remediation. The full equation incorporates AI discovery capability, codebase growth, dependency chain amplification, remediation ceiling, iatrogenic feedback from patches that create new defects, and the unpatchable surface of misconfigurations and architectural design choices.

The Vulnerability Deficit Equation

$$V(t) = V(t-1) + D(t, C(t)) - R\_eff(t) + f(R(t)) + M(t)$$

Every force driving discovery is compounding. AI capability is an emergent property of general AI improvement, not a frontier-lab artifact that will be contained. The global codebase is expanding at an accelerating rate. Dependency chains amplify every vulnerability across thousands of applications. Meanwhile, every force constraining remediation is linear. The CISA/Qualys data, 1.1 billion remediation records across ten thousand organizations, proves the ceiling empirically: a 6.5x increase in effort produced worse outcomes.

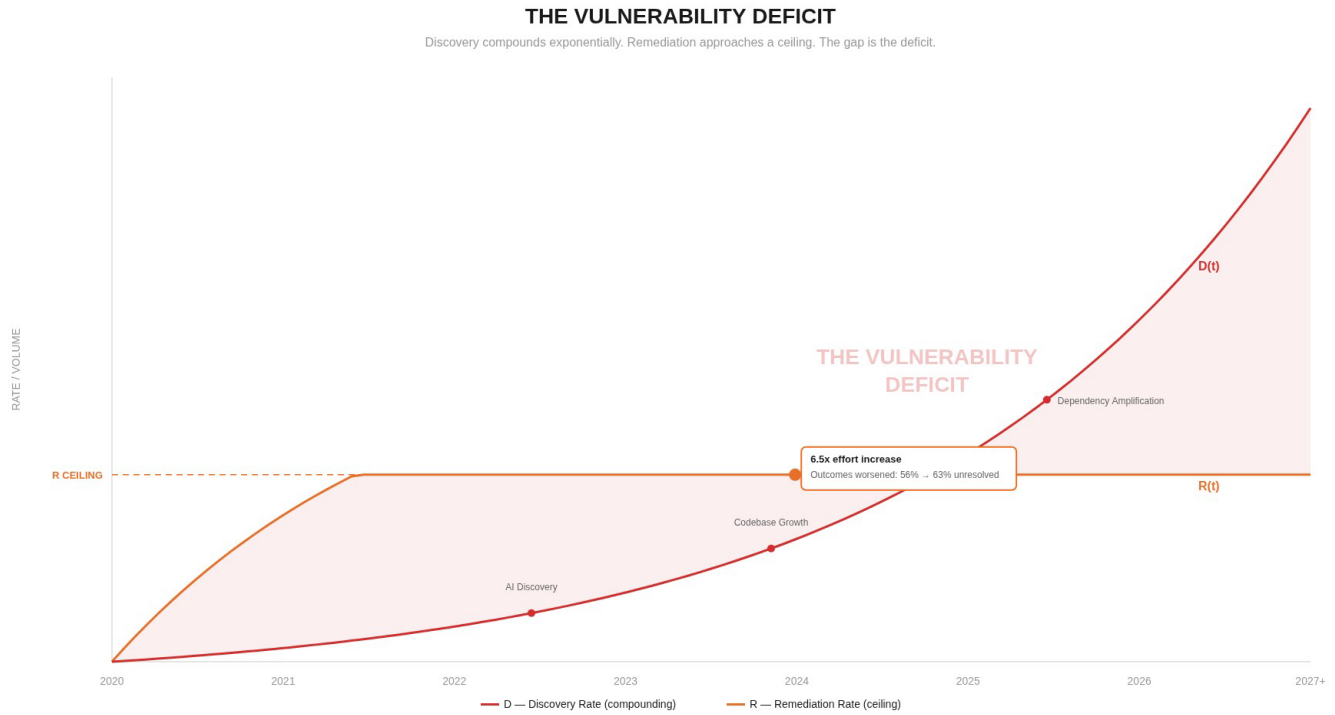

Figure 2: The Vulnerability Deficit

The structural asymmetry is the deeper reason the math cannot be rescued. The attacker needs to find and exploit one vulnerability. The defender relying on remediation needs to find and patch all consequential vulnerabilities before they are exploited. As V diverges, the defender's burden grows without bound while the attacker's remains unchanged.

This asymmetry is specific to remediation as a strategy. A defense strategy based on containment changes the asymmetry entirely. If a compromised workload's communication policy permits zero lateral movement, the attacker's options collapse from thousands to zero, whether or not the vulnerability that enabled the initial compromise has been remediated. The blast radius is one workload.

Containment does not reduce V. It makes V less consequential. It changes the question from "can we patch everything" to "when something is exploited, how far can the damage spread."

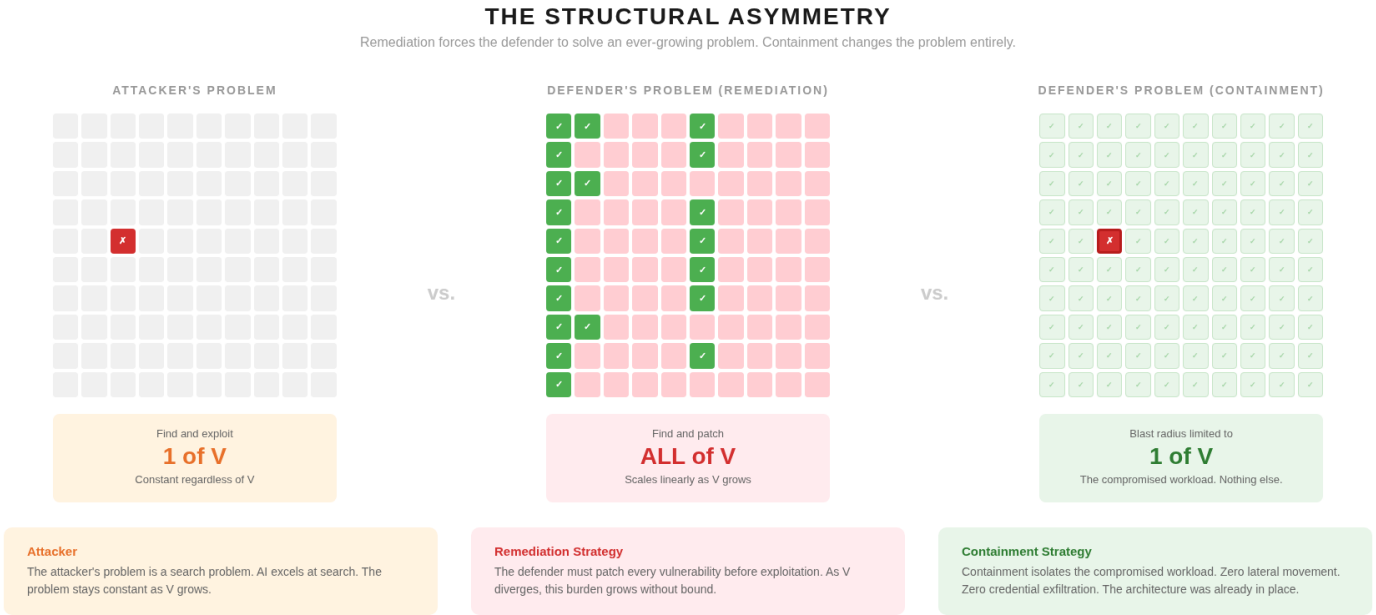

Figure 3: The Structural Asymmetry

Vulnerabilities Are Not the Dominant Vector

The Vulnerability Deficit Equation describes a divergent system inside the domain of vulnerabilities. The deeper problem is that vulnerabilities are not even the dominant attack vector.

82% of intrusions in 2026 ride valid credentials through legitimate channels. Compromised credentials. Rogue employees. Negligent employees. SaaS supply chain trust handed to an attacker by a third party your vendor uses. Vulnerability management cannot reach any of this. There is no scanner for a phished employee, an insider recruited over Telegram, or a Context.AI to Vercel OAuth pivot. The vector that decides most breaches is structurally outside the scope of patching, scanning, and detection. The 2026 breach landscape makes this concrete. The Vercel breach moved through a compromised AI productivity tool with a legitimate Google Workspace session. The Match Group breach rode a phished Okta SSO credential into the AppsFlyer mobile marketing platform. The Crunchyroll breach rode a single phished Telus contractor credential across the trust boundary into Crunchyroll's Slack, Zendesk, and Google Workspace tenants. The McGraw-Hill breach exploited a Salesforce misconfiguration to enumerate 45 million records. None of these required a vulnerability. All of them rode valid credentials through legitimate channels.

A defense strategy that prioritizes patching is optimizing for a vector that decides 18% of intrusions, while leaving the 82% vector structurally unaddressed. Containment is the only control that holds equally against both. It does not care whether the workload was compromised through a vulnerability or a credential. It cares what the workload

The Exploitation Window Has Collapsed

The Vulnerability Deficit equation is generous to remediation, because it assumes defenders learn about vulnerabilities in time to act. CISA KEV data shows a median time-to-exploit of negative seven days for actively exploited vulnerabilities. Attackers are exploiting these vulnerabilities a full week before they appear in any public database. You cannot patch what you do not know exists.

The trajectory shows exponential decay. Time-to-exploit collapsed from 771 days in 2018 to 6 days in 2023, to 4 hours in 2024, to under one day in 2026. The foundational sequence of vulnerability management, discover, disclose, patch, deploy, is broken at step one for a growing class of threats.

When the exploitation window is negative, remediation is not a defense strategy. It is a cleanup activity. The only control that operates in negative time is architecture that was already in place.

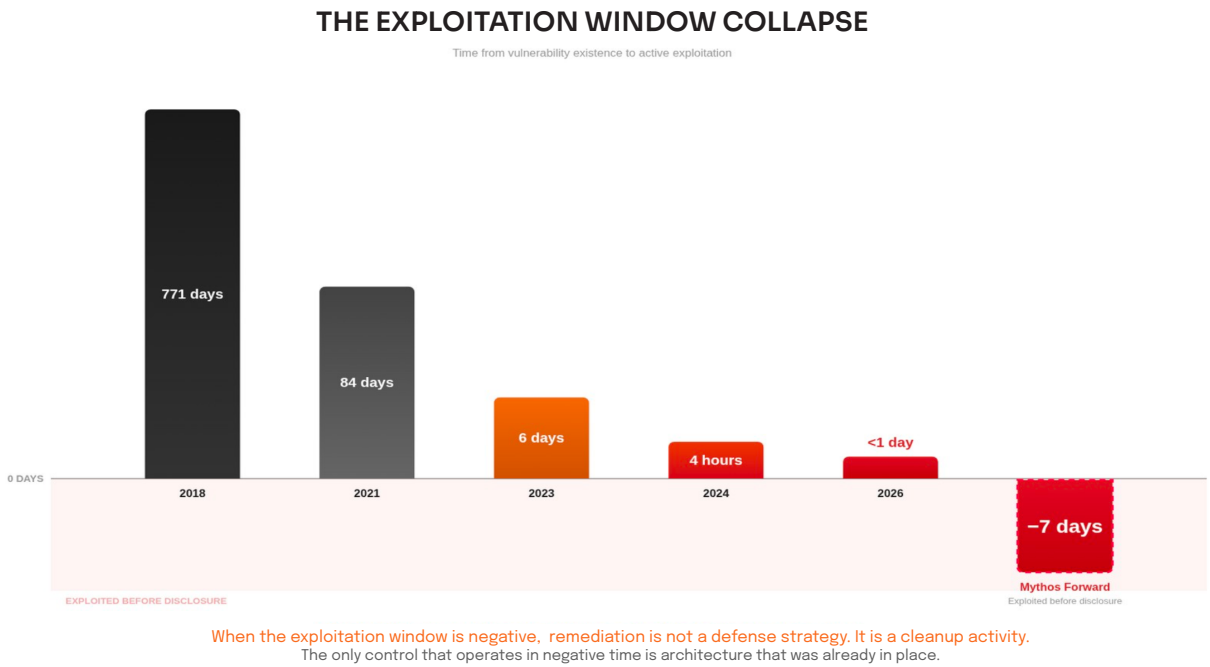

Figure 4: The Exploitation Window Collapse

The Priority Inversion

The Mythos report rates Continuous Patching as CRITICAL and Hardening as HIGH. This ordering rests on an implicit assumption: that increasing remediation velocity is the highest-leverage response to AI-accelerated threats. The Vulnerability Deficit equation proves this assumption is structurally incorrect. The credential vector evidence proves it is incorrect even before the math is applied.

If discovery is compounding on multiple independent curves while remediation has a demonstrated ceiling, and if the dominant attack vector is outside the scope of remediation entirely, then the primary defense variable is not how fast you patch. It is how much damage an unpatched vulnerability or an abused credential can cause. That is a containment question, not a remediation question.

The Mythos report's own case study proves the point. The report describes the LiteLLM compromise in detail: TeamPCP published a malicious package, forty thousand environments pulled it within three hours, and credentials were exfiltrated to attacker-controlled infrastructure. Under the report's own priority ordering, an enterprise would invest first in accelerating its ability to detect and patch the compromised package, a process that at best would take hours to days. The exfiltration took three hours. An enterprise that had prioritized containment architecture, governing every workload communication path, would have had policy in place before the compromise occurred. As documented in Paper 3, one Fortune Global 500 enterprise stopped the exfiltration in exactly this way, not because they detected the compromise, but because the architecture did not permit that workload to reach that destination.

The SANS report's own case study proves that containment, not patching, was the control that determined the outcome.

The Eight Axioms

The Vulnerability Deficit equation describes a system in divergence. The following eight axioms establish why containment is the only architecturally sound response. Each axiom is a structural truth that holds regardless of vendor, implementation, or market category.

Axiom 1: Trust will be violated in any system of sufficient complexity.

In any complex system with enough components, actors, and interdependencies, some element will behave in a way it was not intended to behave. A credential will be stolen. A dependency will be compromised. The question is never whether trust will be violated. It is what the system permits when trust is violated.

Axiom 2: Detection requires distinguishability.

If the attack is the expected behavior, valid credentials, trusted channels, signed packages, detection has no signal to find. You cannot distinguish what is indistinguishable. This is a boundary condition, not a capability gap. The 82% credential vector is the practical proof.

Axiom 3: A centralized inspection point can only govern traffic that traverses it.

Kubernetes pod egress, serverless function communication, VPC peering, hidden communication paths: these are standard cloud networking constructs. The traffic exists. The inspection point does not see it.

Axiom 4: Cloud infrastructure is permissive and internet-addressable by default.

Any workload with an endpoint is accessible to eight billion people. Security groups allow broad access. Communication Governance requires deliberate, active configuration. The default state is open. The shared responsibility model places interior security on the tenant. Most tenants have not built the interior walls. This is not a tenant failure. It is a structural condition of the cloud the providers built.

Axiom 5: AI offensive capability is emergent, permanent, and democratized.

Anthropic did not train Mythos for vulnerability discovery. The capability emerged from improvements in coding, reasoning, and autonomy. Mythos is one data point on a curve, not the curve itself. Open-weight models with full capabilities are available under permissive licenses. There is no mechanism to recall or restrict them. The offensive capability they enable is permanent, and the cost of wielding it is collapsing toward consumer compute.

Axiom 6: Containment is architecturally independent of detection.

A wall between two rooms works regardless of whether you know someone is in one of them. Containment does not require detecting the compromise first. The wall exists before the breach.

Axiom 7: Ephemeral workloads cannot be secured by models that assume persistence.

A container that exists for sixty seconds cannot be protected by a security model that requires agent installation, identity enrollment, and human review. The security model takes longer than the workload's life. AI workloads are the apex of this problem, because they are the most ephemeral, most privileged, and most rapidly shipped class of workload in the enterprise.

Axiom 8: Attack surface exploitation scales with the number of capable attackers.

The surface did not change. The number of hands testing every door did. AI changed this by orders of magnitude. Industrialization changed it by another order of magnitude on top.

These axioms are not vendor claims. They are structural truths drawn from networking physics, cloud provider documentation, empirical measurement, and information theory. The chain of reasoning they produce is direct: trust will be violated (1). For a growing class of attacks, the violation will be indistinguishable from legitimate activity (2). Centralized inspection cannot see traffic that does not traverse it (3). Cloud is permissive by default and the tenant is responsible for interior security (4). AI is making offensive capability emergent, permanent, and democratized (5). The only architecturally sound response is containment that is independent of detection (6), that can secure workloads regardless of lifespan (7), and that holds as the number of capable attackers increases by orders of magnitude (8).

The Honest Boundary

Intellectual honesty requires distinguishing between what is proven and what is predicted. The structural claims in this paper, the Vulnerability Deficit equation, the measured remediation ceiling, the 82% credential vector data, the architectural independence of containment from detection, and the structural truth of each axiom, hold regardless of market conditions or organizational choices. The argument for re-prioritizing PA 8 above PA 5 rests entirely on these structural claims. It does not depend on any prediction, any vendor capability, or any timeline estimate. The math holds regardless.

The Question Before the Industry

The SANS Mythos Report is the most authoritative document the cybersecurity industry has produced on AI-accelerated threats. This paper does not challenge its authors' expertise. It challenges a single prioritization decision, one that the evidence, including the report's own evidence, does not support.

Continuous patching is necessary. It will always be necessary. Patching is hygiene. But hygiene is not architecture. And when the exploitation window has collapsed, when discovery compounds on curves that remediation cannot match, when 82% of intrusions use valid credentials that no scanner will find, and when the toxic combination is concentrating its damage on AI workloads that traditional controls cannot reach, the highest-priority action is not to patch faster. It is to ensure that when the next compromise occurs, the blast radius is determined by architecture, not by luck.

Three questions can clarify where your organization stands:

Does your enforcement govern every path a workload can take to any destination, including Kubernetes pods, serverless functions, east-west traffic, and regions where policy has not yet propagated?

If a compromised workload attempted to exfiltrate credentials right now, would your architecture stop it, or would you learn about it from a log you review tomorrow?

If a valid credential is used against one of your AI workloads at 3:00 AM on a Sunday, what, architecturally, prevents it from reaching the ledger?

If the answer to any of these requires a qualification, then PA 8 is your highest priority. Not because this paper says so. Because the math does.

The SANS Mythos Report identified the right actions. The Vulnerability Deficit equation, the credential vector, and the toxic combination reveal the right order. Containment is not the second priority. It is the foundation on which every other priority depends.

The axioms in this paper describe the architectural properties any containment solution must have. Aviatrix built it. The Cloud Native Security Fabric governs every workload communication path, across every cloud, every VPC, every Kubernetes cluster, every serverless function, from a single policy plane. One rule. Universal propagation. Enforced at the workload, not at a chokepoint. The Fortune Global 500 enterprise that stopped the LiteLLM exfiltration did so because that architecture was already running.

This is Paper 4 of 4 in the Containment Era series.

Paper 1: The Containment Era – Why the Threat Model Outgrew the Architecture.
Paper 2: The Containment Platform – How Cloud Native Security Fabric Closes the Architectural Divide.
Paper 3: 144 to 1 – Why Every Workload in Your Cloud Is Already Exposed.