The “AI Vulnerability Storm”: Building a “Mythos-ready” Security Program

Expedited Strategy Briefing
By the CSA CISO Community, SANS, [un]prompted,
the OWASP Gen AI Security Project, and the wider community.

Contact cisos@cloudsecurityalliance.org with any inquiries.

16 April, 2026

Version 0.92
The latest version of this document can be found here.

CSA cloud security alliance® SANS [un]prompted OWASP GenAI SECURITY PROJECT

Authors

Gadi Evron

CEO, **Knostic**, CISO-in-Residence for AI, **Cloud Security Alliance**

Robert T. Lee

Chief AI Officer, Chief of Research, **SANS** Institute

Rich Mogull

Chief Analyst, **Cloud Security Alliance**

Contributing Authors

Jen Easterly

CEO, **RSAC** and Former Director, CISA

Chris Inglis

Former National Cyber Director, **The White House**

Heather Adkins

CISO, **Google**

Sounil Yu

CTO, **Knostic**, former Chief Security Scientist, Bank of America

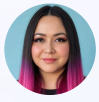

Katie Moussouris

Founder and CEO, **Luta Security**

James Lyne

CEO, **SANS**

Maxim Kovalsky

Managing Director, AI Security CoE, **Consortium Networks**

Joshua Saxe

CTO and Co-founder at **Security Superintelligence Labs**, former AI and Llama Security Lead, Meta

Bruce Schneier

Chief of Security Architecture, **Inrupt**, Fellow and lecturer, Harvard Kennedy School

Phil Venables

Ballistic Ventures, former CISO, Google Cloud

Rob Joyce

Former Cybersecurity Director, **NSA**

Jim Reavis

CEO, **Cloud Security Alliance**

John N. Stewart

Talons Ventures, former CSTO, Cisco Systems

Dave Lewis

Global Advisory CISO, **1Password**

John Yeoh

CSO, **Cloud Security Alliance**

Ramy Houssaini

CCSO, **Cloudflare**

Reviewers

Many CISOs, and some other practitioners, assisted in reviewing and editing this document. These are the ones who would share their names publicly, in alphabetical order (by last name):

Mark Aklian, Founder & CISO, Silver Oak Cyber
Barry Anderson, Information Security Architecture Strategy and Engineering Manager, HESTA
David Aronchick, Co-Founder, Expanso / Kubeflow
Jake Bernardes, CISO, Anecdotes
Alan Berry, CISO, Centene Corporation
Bryson Bort, CEO, SCYTHE
Jeff Bryner, CISO, Independent
Michael Calderin, CISO
Daniele Catteddu, Chief Technology Officer, Cloud Security Alliance
Viswanath Chirravuri, Global Product Security Director, Thales
Mea Clift, CISO, Cengage
Scott Clinton, Co-chair, OWASP Gen-AI Security Project
David B. Cross, CISO, Atlassian
Chris Cochran, Field CISr,O & VP AI Security, SANS Institute
Michael Colao, former Corporate CISO, AXA, Director, Island Cyber
Daniel Cuthbert, Associate Fellow, Cyber and Tech, RUSI
Julie Davila, VP Product Security, GitLab
Michael Douglas, Managing Partner InfoSec Innovations, and SANS Instructor, SANS Institute
Yoni Efrati, former CISO, Bank HaPoalim
Eliya Elon, EIR, Notable Capital
Sergej Epp, CISO, Sysdig
Chris Farris, Cloud Security Nerd, Securosis
Alex Foley, Data Security TISO, Wells Fargo
George Gerchow, CSO, Bedrock Data
Dan Glass, CISO, Delek US Holdings
Barry Greene, Co-Founder, Qubit Cyber
The grugq, Independent Security Researcher, Independent
Erik Hart, Global CISO, Cushman & Wakefield
Gary Hayslip, CISO in Residence, Halcyon, former CISO, SoftBank
Dustin Heywood, Senior Technical Staff Member, IBM
Ariel Herbert-Voss, CEO, RunSybil
Heather Hinton, CISO, Sitecore, Lecturer, Harvard Extension School
Dave Hoelzer, independent
Matt Holland, Director of Cyber and AI Security, Kainos
Igor Ignatov, Head of Security & Compliance, Cognichip Inc.
Waylon Janowiak, Head of Information Security & IT, Faire
Mike Johnson, CISO, Rivian
Avner Langut
Paul Lanzi, Principal Consultant, IDenovate
Rock Lambros, Director of AI Standards and Governance, Zenity, Founder, RockCyber
Josh Lemos, CISO, lululemon
Peter Liebert V, CISO, Salesloft
Ariel Litvin, former CISO, First Quality Enterprises
Bob Lord, former CISO, Yahoo, CSO, DNC
Myke Lyons, CISO, Cribl
Ciaran Martin, Head of Cyber Leaders Network, SANS Institute, Founder and former CEO, UK NCSC
Michael Machado, CISO & CDO, Hyland
Donald McFarlane, Principal Technical Advisor, Microsoft
Gal Malach, CTO and Co-Founder, Terra Security
Tomas Maldonado, CISO, NFL
Greg McCord, Founder and CISO, McCord Keystone Advisory
Ross McKerchar, CISO, Sophos
Greg Notch, CSO, Expel
Charles Nwatu, former Head of Security, Netflix
Helen Oakley, OWASP GenAI Security Project
Mark Orsi, CEO, Global Resilience Federation
Teju Oyewole, CISO, Brave Technology
David Quisenberry, CTO & CISO, Ferguson Wellman Capital Management
Rupa Parameswaran, VP, Security and IT, Ex-Handshake, Amplitude
Gavin Reid, CISO, Human Security
Gerardo Richarte, CISO, Satellogic
Joshua Scott, VP of Security and CISO, Hydrolix
Mark Seiden, Volunteer and Security Advisor, Internet Archive
James Shank, Director of Threat Operations, Expel
Samir Sherif, Global Field CISO, Fastly
Conor Sherman, CISO in Residence, Sysdig
Ed Skoudis, President, SANS Technology Institute
John Sotiropoulos, OWASP GenAI Agentic Security Initiative Co-Lead, Co-Founder, Deep Cyber
Sean Todd, CISO, trycoral.ai
Anna Sarnek, Senior Fellow, McCrary Institute for Cyber & Critical Infrastructure Security
Holger Spohn, CISO, Candriam
Rob van der Veer, OWASP AI Exchange
Mikael Vinding, CISO, AP Technology
Yabing Wang, CISO & CIO, Justworks
Mike Wilkes, Adjunct Professor, New York University
Jeff Williams, CISO, Sigma360
Steve Wilson, Chief AI Officer, Exabeam, Co-Chair OWASP GenAI Security Project
Jason Woloz, CISO, TransUnion

All the listed authors and reviewers represent only themselves, and not their employer(s).

Join the Cloud Security Alliance CISO community, email us at cisos@cloudsecurityalliance.org.

Table of Contents

I. Executive Summary

What happened:

- AI, as demonstrated by Anthropic's Mythos, has significantly increased the likelihood of attackers discovering new vulnerabilities, creating new exploits, and using them in complex automated attacks at scale.
- While AI also increases the speed of patch development and reduces defects in new software, defenders still face a heavier relative burden due to the inherent limitations of patching. Attackers gain asymmetric benefits.

How is this different from the status quo?

- In the near term, security organizations will likely be overwhelmed by the need to apply patches and respond to AI-discovered vulnerabilities, exploits, and autonomous attacks.

What to do now to deal with the current risk spike?

- Adjust risk calculations and re-orient security program resources for increasing volume of patches, decreasing time to patch, and more persistent and complex attacks.
- Focus on the basics and harden your environment further. Segmentation, egress filtering, multifactor authentication, and defense-in-depth/breadth all increase the difficulty for attackers.

What do we believe will happen next?

- The storm of vulnerability disclosures from Project Glasswing is the first of many large waves of AI-discovered vulnerabilities that may occur in rapid sequence.
- The capabilities seen in Mythos will quickly become more widely available, dramatically increasing the number and frequency of complex, novel attacks organizations will face.

What else should start now to be ready for the next waves?

- Prioritize robust dependency management to reduce vulnerabilities in third-party and open-source components.
- Enforce automated security assessments consistently in your development processes, including using LLM-powered agents to find vulnerabilities before attackers do.
- Introduce AI agents to the cyber workforce across the board, enabling defenders to match attackers' speed and begin closing the gap.
- Re-evaluate your risk tolerance for operational downtime caused by vulnerability remediation, to account for shorter adversary timelines.
- Update governance for more efficient vendor onboarding and increase headcount to facilitate a faster cycle deployment of new AI-based defenses.
- As an industry we need to strengthen our coalitions, cooperation, and coordination.

II. Key Takeaways for the CISO

Use LLM-based vulnerability discovery and remediation capabilities.

Unlike defensive AI technologies, LLM-based vulnerability discovery capabilities are already mature and can be used to your advantage. Start immediately by asking an agent for a security review of any code, and build toward a VulnOps capability.

Update risk metrics.

With the shifting landscape, many of your metrics and risk assessments may be outdated and could affect business reporting. Consider how to update these, and communicate the challenge with stakeholders.

Accelerate your team by the use of coding agents.

While defensive AI technologies are lagging behind offensive ones, agents can already accelerate human action across the board, from incident response to GRC. Encourage and require your team to use these agents to accelerate their capabilities. Triage and test patches, red team your environment, automate audit data collection, and accelerate security operations overall.

Prepare to respond to more incidents.

Run tabletop exercises for multiple simultaneous high-severity incidents occurring within the same week, and have playbooks in place for high-level, critical incidents. Examine how to automate remediation capabilities to the degree possible. Verify and enable mitigating controls such as segmentation, egress filtering, Zero Trust architectures, phishing-resistant MFA, and secrets rotation to limit impact when exploitation occurs. The supply chain will be affected.

Increase focus on the basics.

The basics remain valid and can be prioritized for risks that cannot otherwise be mitigated. Segmentation, patching known vulns, Identity and Access Management, and defense-in-depth/breadth all increase the difficulty for attackers. To lower latent risk, expanding these efforts while there is time is prudent.

We cannot outwork machine-speed threats. Re-prioritize, automate, and prepare for burnout.

The cadence and volume of vulnerability disclosures will exceed anything we have experienced before. Consider how you manage current priorities, and request additional headcount and budget for reserve

capacity to avoid exhausting available resources, or potentially burning out existing staff. This, in parallel with adoption of coding agents, re-prioritization, putting more automation in place, and helping your team through career uncertainties and upskilling challenges.

Evolve to a Mythos-ready Security Program.

Mythos is likely one of many changes coming to cybersecurity risk. If not already underway, seriously consider incorporating Mythos and its implications into your strategy.

Build Collective Defense Now.

Attackers already operate as syndicates, crowdsourcing, sharing tools, and moving as a collective. Engage now with sector coordinating groups, ISACs, CERTs, and standards bodies to share threat intelligence, coordinate response, and produce sector-specific guidance for this moment. Defenders must do the same and leverage our coordinating groups, especially when considering organizations that fall below the Cyber Poverty Line, as introduced by Wendy Nather.

III. Introduction

Many of our assumptions about the capabilities of AI in vulnerability research, exploitation, and autonomous attacks may be outdated. Throughout 2025 and into 2026, we've seen continuous examples of increasing capabilities, both in research and in actual in-the-wild attacks. AI-driven vulnerability discovery and exploitation has been accelerating for over a year. See Appendix A for more details and historical evidence.

Anthropic's Claude Mythos (Preview) represents a step change in that trajectory, autonomously finding thousands of critical vulnerabilities across every major operating system and browser, generating working exploits without human guidance, and empowering autonomous attack orchestration, all at a speed and scale that outpaces any prior capability.

The asymmetry this creates is structural. AI lowers the cost and skill floor for discovering and exploiting vulnerabilities faster than organizations can patch them. The window between discovery and weaponization has collapsed to hours. Attackers gain disproportionate benefit, and current patch cycles, response processes, and risk metrics were not built for this environment.

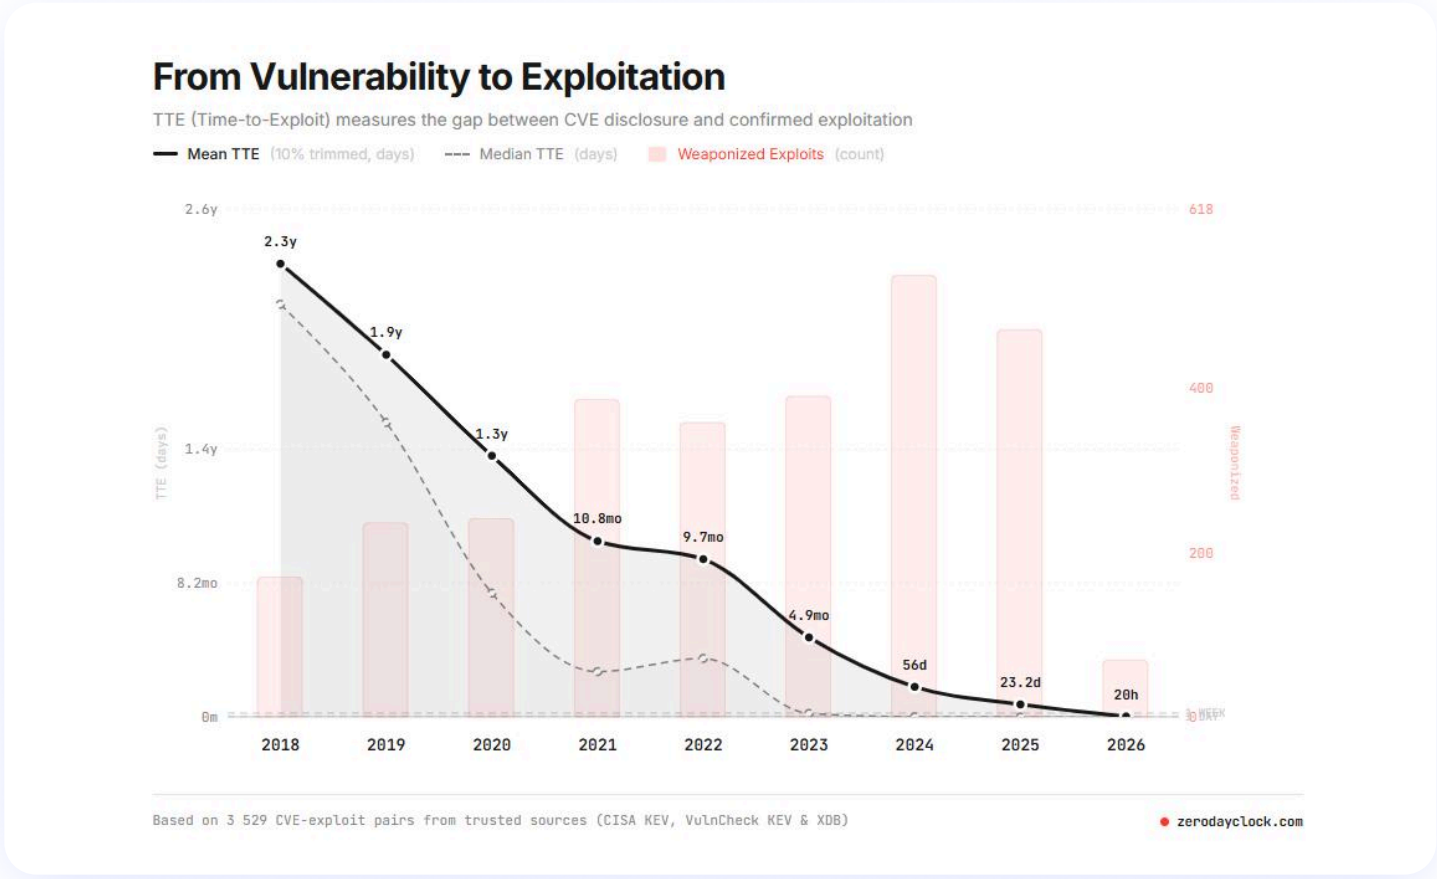

Diagram from the Zero Day Clock, by Sergej Epp, demonstrating the collapsing time to exploitation, which is now down to hours.

While many of these capabilities pre-date this new model, Mythos-class capabilities do represent a step-change, and will proliferate. The organizations that respond well will be those that build the muscle now: the processes, the tooling, and a culture willing to adopt AI as a core part of how security gets done. That adaptability will help determine who meets the next wave on their own terms.

This moment requires reprioritizing resources, reviewing risk levels and controls, and leveraging AI where feasible. At the time of this writing, most AI defensive controls and approaches are not yet mature. That said, AI attacker technology may be used for defense purposes and coding agents will help.

The detailed recommendations are included later in this document.

Mythos & Glasswing: Why They Matter

Mythos

Mythos is distinguished from previous capabilities on both technological and strategic levels, even if many of its attributes **already existed** and have evolved over the past year. Technologically, models like Mythos exhibit three capabilities that make them different:

1. **Exploits without scaffolding.** Internal lab environment testing at Anthropic showed Mythos generated 181 working exploits on Firefox where Claude Opus 4.6 succeeded only twice under the same conditions, marking a substantial jump in autonomy and reliability.
2. **Complex, chained vulnerabilities.** Mythos identifies vulnerabilities composed of multiple primitives chained together, such as scenarios requiring multiple memory corruption bugs combined into a single exploit path.
3. **“One-shot” (single-prompt) capability.** Mythos accomplishes significantly more with a single prompt, without elaborate scaffolding or agent configuration.

Strategically, Mythos broke into mainstream media beyond technical security communities and reached into boardrooms, raising awareness and the urgency of AI-driven vulnerability risks. This has forced security teams to respond and opened the door for new resources and funding across the industry.

Glasswing

The scale and speed of Mythos prompted Anthropic to create **Project Glasswing**, possibly the largest multi-party vulnerability coordination effort in history. Anthropic provided selected critical infrastructure providers, industry partners, and open source maintainers early access to Mythos so they could patch their own products. Other AI model vendors have launched similar vetted-participant programs.

The most significant limitation of Project Glasswing is that it can only cover so much.

The world’s exploitable attack surface is vastly larger than what any curated partner ecosystem can cover, and most organizations that build or maintain critical software will not have early access to Mythos-class capabilities. Meanwhile, the competitive landscape is narrowing that window. If comparable offensive capabilities emerge in other frontier models within months, and in open-weight models within six months to a year, the defensive advantage conferred by early access becomes time-limited by definition.

While the coordination model Glasswing established is critically important, its impact will depend heavily on how quickly it can expand coverage, and whether the patch and disclosure pipeline can keep pace with both AI progress and adversarial adoption.

The Evolution of LLM-based Offensive Capabilities, 2025/6

Jun 24, 2025

XBOW tops the HackerOne leaderboard

- **XBOW** became **#1 on HackerOne's US leaderboard**, first autonomous system to outperform all human hackers on the platform
- Open-source **raptor** demonstrated that autonomous **vulnerability research** is available to anyone using an off-the-shelf agent

Aug 5, 2025

Google Big Sleep finds 20 real-world zero-days

- Google's **Big Sleep discovered real zero-days in open source**, 20 vulnerabilities in projects including FFmpeg and ImageMagick, each found and reproduced autonomously

Aug 8, 2025

DARPA AIxCC finals at DEF CON 33

- **DARPA AIxCC found 54 vulnerabilities in four hours of compute** across 54 million lines of code

Sep 2025

Singularity warning issued

- Heather Adkins (CISO, Google) & Gadi Evron (CEO, Knostic) **publish warning** that attackers are racing toward a singularity moment
- Autonomous vulnerability discovery and exploitation estimated ~6 months away

Nov 14, 2025

First AI-orchestrated espionage campaign disclosed

- Anthropic **disclosed a Chinese state-sponsored group** had used Claude Code to autonomously run full attack chains — recon through exfiltration — across ~30 global targets (detected mid-Sep 2025)

Feb 5, 2026

AI finds hundreds of high-severity bugs; autonomous attacks discovered

- Anthropic (using Claude Opus 4.6) **reported 500+ high-severity vulnerabilities** in open source software
- AISLE found **12 OpenSSL zero-days**, including a CVSS 9.8 flaw dating to 1998
- Sysdig **documented an AI-based attack** reaching admin-level access in 8 minutes
- Gambit released a report on the **AI-led compromise** of Mexican government infrastructure

Mar 2026

Open source projects overwhelmed

- Linux kernel bug reports climbed from 2 to 10/week, initially hallucinated, now all verified real
- The **curl project**, which **discontinued its bug bounty** over AI-generated "slop" reports, now echoes the same shift, an increasing share of reports are quality AI-supported findings

Mar 2026

[un]prompted conference + Zero Day Clock

- **[un]prompted** introduces multiple talks, open source projects, and specific implementations demonstrating the risk with data
- The **Zero Day Clock** is launched, visualizing the collapse of time-to-exploit — now under one day in 2026

Apr 7, 2026

Claude Mythos Preview

- Anthropic announces Claude Mythos Preview & Project Glasswing. Discovers thousands of zero-days across every major OS and browser. 72% exploit success rate. 27-yr-old OpenBSD bug.

IV. The Mythos-ready Security Program

The changing landscape, and resulting risk and impact, require an approach that is both operational, incident response-like, and strategic, focused on program building over time. This is implemented across three time horizons.

It is beyond the scope of this text to be exhaustive or prescribe how a full-fledged AI security program should be built. Rather, we selected *high-impact* recommendations that you can start with today, based on what the community can clearly discern at this early stage.

Beyond Application Security and Vulnerability Management, Mythos affects the wider security program. For example:

- **Operationally**, expect a potential deluge of new patches released from the 40 vendors in the early access program, similar to recent experience of needing to respond to multiple supply chain incidents within a two-week timeframe.
- **Risk management-wise**, business risk is shifting and engagement with stakeholders on risk planning and tolerance is key. The CISO's ability to manage risk has been reduced to a degree that could potentially have effects on business reporting and projections.
- **Strategically**, longer-term gap analysis and selective overhaul of various functions will be beneficial, including governance processes to support faster technology onboarding and the implementation of innovative AI-based security controls.

To start, a Mythos-ready security program should achieve **minimum viable resilience**. It would upgrade and realign measurements to a higher maturity level on key metrics such as cost of exploitation, early detection of compromise, and blast radius containment.

This matters because many of the assumptions underlying our cyber defense programs are being challenged. For example, time to exploitation has been reduced to minutes, we can no longer assume a patch will be ready in time for remediation purposes, incident frequency is likely to increase, the CVE system may not scale, shadow IT will fragment central control as coding agents proliferate to Citizen Coders, employees develop their own infrastructure, and threat intelligence is lagging behind on vulnerability discovery and exploitation.

The First of Many Waves?

Further, any program we build must acknowledge Mythos is only the first wave of future AI technology disruptions. In building a Mythos-ready program, we are not only seeking a return to equilibrium but also preparing to maintain balance for the waves ahead.

A Mythos-ready program should also account for how these shifts affect your team. The pace of change is real, and practitioners across all levels are working through what AI means for their roles and skills. This is a normal response to disruptive capability shifts, not a crisis of relevance. The practitioners who adapt fastest will be the ones who lean into AI

tooling rather than viewing it as a threat to their expertise.
The path forward is doubling down on fundamental security controls and hands-on adoption of agents at every level, from the CISO down. Every security role is becoming an "AI builder" role, and the barrier is lower than most people realize. Using a coding agent is now **easier than using Excel.**

Are We Outmoded? The Human Cost and the Opportunity Ahead

Leaders must be clear-eyed about the human cost of this transition. This isn't just about burnout, but rather about reprioritization, automation, and an opportunity for clarity in communications and personal growth.

Security teams are caught in a vice: AI is simultaneously accelerating the frequency of of vulnerability reports they must respond to, the volume of code their organizations are shipping, and the expanding attack surface.

Beyond a workforce already at capacity, absorbing exponential increases in workload, staff is operating under increased uncertainty.

They often feel they are falling behind from a skills perspective, are concerned about being replaced by AI, fall while handling the cognitive intensity of management demands of integrating AI into their own workflows. Often, without reprioritization from management. corresponding investment in automation and tooling, or headcount and well-being.

- Burnout and attrition in security functions represent a direct operational risk - the expertise needed to navigate this transition is scarce, takes years to develop, and is not replaceable on short timescales. Security team resilience, including sustainable workload, mental health support, and retention, should be treated as a strategic priority with the same urgency as the technical challenges AI presents.

- Security practitioners, ourselves included, are facing a culture challenge. Many are uncertain about how their roles will evolve. It is often unclear to them, and us, how we could keep up with the pace of change. This affects even the most technical, such as vulnerability researchers, many of whom are asking questions about the future and if they will have a place in it.

For now, we are not outmoded. Agents, often in the form of coding agents, also represent an opportunity for personal growth, and a feeling of empowerment. Everyone on your team, including you, can become hands on. All roles will likely become "AI builders", where technical skills and specific domain know-how are augmented by agents. Getting started is now easier than using Excel. All you need to know is English.

The Shrinking Time Horizon

The time available for action is shrinking, and we need to find ways to move faster. Long-term goals should be considered a quarter away, at most.

10 Questions to Understand Your Security Program State and Influence

A questions-based approach to triage your understanding of your security program, to reach ground truth, as well as gauge your influence on various business functions.

Question	Context
What is our actual stance on AI today?	Allowed, tolerated, restricted, or unknown.
Can employees use agentic coding tools in the enterprise today?	Making use of agentic capabilities such as looping LLM tool use, and specifically coding agents (regardless of writing code), not just chatbot access. Do you have security guardrails in place for these coding agents?
Can employees contribute to open source without legal ambiguity?	A legal and IP question, not a technology philosophy question.
Do we have disciplined control repos, artifacts, and software, including for agentic supply chain such as MCP servers, plugins, and skills?	Source control, package paths, artifact provenance, and what is actually allowed in, in the CI/CD pipeline and through coding agents.
Is there a real cooling-off point/security gate between code change and production?	Demonstrates enforcement of security in release cycles and control of software supply chain.
Is security operational, or primarily advisory?	The extent to which the security function can directly affect outcomes, or does it serve mostly as a review and escalation function.
What is the fastest this company has made a security-driven production change in the last year?	Use a real example, not a policy statement.
Are our critical "crown jewels" explicitly tracked and current?	Not theoretically important systems. The actual few that matter most, and their main dependencies.
Do we know how to get urgent work prioritized by our key third parties?	Feature requests, bug reports, security escalations, relationship ownership, and leverage.
Does executive leadership have a working definition of urgency?	If everything is a crisis, nothing is urgent.

Updating Your Security Program

With those answers, we start with a draft risk register, followed by a list of prioritized actions and controls for a Mythos-ready security program, based on what the writers believe are most likely to be effective and impactful for most organizations.

Then, provide with an action plan for consideration in updating your security program.

Each action below is broken down into when it should commence, and a generalized estimate on a potential risk is linked to recognized frameworks, and each action to a time horizon under which it could be completed, for most organizations.

See appendix B for a full legend.

A Mythos-ready Security Program Risk Register (DRAFT)

	Severity	Risk	Description	Type	Framework Refs	Maps to Priority Action
			CRITICAL			
1	CRITICAL ⚠⚠⚠	**Accelerated Threat Exploitation** AI-autonomous exploit generation at machine speed	AI models have been discovering vulnerabilities and creating exploits for over a year. Mythos accelerates this significantly, but the capability predates it. What changes is the speed, scale, and the reduction in skill required to execute complex attacks, democratizing capabilities that were previously expensive and skill-intensive. Non-frontier, open-weight models can already achieve much of this at an accessible cost. Frontier models like Mythos are the acceleration, not the starting gun. Each patch also becomes an exploit blueprint, as AI accelerates patch-diffing and reverse engineering of fixes.	Threat	AML.T0040, AML.T0043, PR.PS, PR.IR	**PA 4, 5**
2	CRITICAL ⚠⚠⚠	**Insufficient AI Automation Capabilities** Defenders operating at human speed while attackers operate with AI augmentation	Attackers freely use AI coding agents for vulnerability discovery, exploit development, and attack orchestration. Many defensive teams are not yet aware of equivalent capabilities available to them, or lack the security controls to deploy them confidently. The resulting asymmetry is not just technological but cultural: teams that do not adopt AI agents cannot match the speed or scale of AI-augmented threats, regardless of their technical skill.	Capability gap	GV.OC, GV.RM, DE.CM, RS.MA	**PA 1, 2**
3	CRITICAL ⚠⚠⚠	**Unmanaged AI Agent Attack Surface** Privileged AI agents outside existing control frameworks	Agents are necessary to counter AI-speed threats, but they are privileged, insecure by default, and not covered by existing security controls. This asset class introduces defensive risks (insecure, privileged agents within your own	Vulnerability	LLM06, ASI02, ASI03, AML.T0047, PR.AA, GV.SC	**PA 3**

			environment) and supply chain risks (compromised or manipulated agents from third parties). These have different owners and different mitigations.			
4	CRITICAL	**Inadequate Incident Detection and Response Velocity** Detection and response at human speed against machine-speed attacks	AI has reduced the sophistication and time needed to construct complex attacks. Defensive detection and response capabilities have not yet been upgraded to match, creating an asymmetric speed advantage for attackers. Alert triage volumes, SIEM correlation speed, and containment authorization latency were designed for human-paced threats.	Capability gap	ASI08, AML.T0047, DE.CM, DE.AE, RS.MA	**PA 9, 10**
5	CRITICAL	**Cybersecurity Risk Model Outdated** Stakeholder decisions based on pre-AI risk models	Security reporting metrics built on pre-AI assumptions about exploit timelines and attack complexity may no longer reflect actual exposure. The CISO's ability to control risk has shifted, which could affect business reporting and projections. Outdated risk models could lead to underfunding of the controls that prevent incidents.	Governance	GV.OC, GV.RM, RS.CO	**PA 6**
			HIGH			
6	HIGH	**Incomplete Asset and Exposure Inventory** Unknown attack surface, assets, code, dependencies, shadow agents	AI-accelerated attacker capabilities change which assets are at highest risk and which controls matter most. Attackers can now scan an entire OS codebase at accessible cost and enumerate your exposure faster than you can inventory it. For assets that cannot be patched or directly defended, inventory determines whether you can segment, isolate, or monitor them. Without continuously updated inventory, controls have inherent gaps. The proliferation of coding agents to non-developer users further fragments central IT visibility.	Vulnerability	ASI04, AML.T0000, ID.AM, GV.SC	**PA 7**
7	HIGH	**Unsecured Software Delivery Pipeline** Code shipping without AI-driven security review	Code produced by both humans and AI agents ships without consistent security review. AI-generated code introduces vulnerabilities at higher volume than manual development. The risk compounds: more code produced faster, with the same defect rate, against a more capable adversary. Without LLM-driven review integrated into the pipeline, exploitable flaws reach production before defenders can find them.	Vulnerability	LLM01, LLM05, LLM08, ASI01, AML.T0018, AML.T0051.001, PR.PS, ID.IM	**PA 1**
8	HIGH	**Network Architecture Insufficient for Lateral Movement Containment**	A flat or insufficiently segmented network gives every successful exploit leverage. AI-driven attacks worsen this: automated multi-hop lateral movement exploits poor architecture faster and more creatively than manual attackers ever could.	Vulnerability	PR.IR, PR.PS	**PA 8**

		Flat or insufficiently segmented network enabling 1:N exploit leverage	When AI-accelerated vulnerability discovery increases the volume of exploitable findings, architectural segmentation becomes the primary control limiting blast radius.			
9	HIGH	**Continuous Vulnerability Management Maturity Gap** Reactive posture against continuous AI-discovered zero-days, no VulnOps function	AI-driven vulnerability discovery, which predates Mythos but is significantly accelerated by it, means zero-day vulnerabilities in your own code and third-party software can be discovered and weaponized before your security team knows they exist. Quarterly pen tests and reactive patching cycles cannot keep pace with continuous AI-driven discovery. Existing CVE/NVD infrastructure and patch prioritization workflows were built for dozens of critical CVEs per month, not hundreds.	Capability gap	ASI10, ASI06, AML.T0018, ID.RA, ID.AM, DE.CM	**PA 11**
10	HIGH	**Threat Detection Dependent on Lagging Intelligence** CVE- and KEV-based intelligence structurally outpaced by AI discovery rates	Threat intelligence has been falling behind AI-accelerated vulnerability discovery for over a year. Mythos widens the gap further. The CVE system may not scale to AI-generated discovery rates, and novel vulnerabilities have no listing in KEV by definition.	Capability gap	AML.T0000, DE.CM, ID.RA, GV.OV	**PA 9, 10**
11	HIGH	**Innovation Governance and Oversight Deficit** Governance vacuum creating approval friction that slows defensive AI adoption	Without a cross-functional governance mechanism, the onboarding and deployment of any new control runs into approval friction that slows adoption. This is where the liability and governance asymmetry gets addressed structurally. AI-accelerated timelines mean this friction now has a harder deadline.	Governance	GV.OC, GV.RM, GV.RR, GV.OV	**PA 2, 4**
12	HIGH	**Regulatory and Liability Exposure from AI-Discovered** Vulnerabilities Shifting standard of care as AI scanning becomes broadly available	The EU AI Act (August 2026) introduces automated audit, incident reporting, and cybersecurity requirements around AI. Existing regulations use reasonableness as a test. When AI can find significantly more vulnerabilities at accessible cost, the standard of what constitutes reasonable defensive effort shifts. Boards will face questions about whether they used available AI tools for defensive scanning, and whether not doing so constitutes negligence. This is a governance risk with direct financial exposure.	Governance	GV.OC, GV.RM, GV.RR	**PA 1, 4**
			MEDIUM			
13	MEDIUM	**AI Hype and Confusion Causing Systematic Inaction** Signal-to-noise collapse in threat and technology guidance	The volume of AI-related security guidance, commentary and vendor claims exceeds anything the industry has experienced. Security leaders find it difficult to navigate the noise. The confusion itself is a consequential risk: teams that dismiss the shift as hype, or exhaust their attention on low-signal content, will miss critical threat landscape changes they need to react to.	Governance	GV.OC, GV.RM	**PA 1**

Priority Actions for a Mythos-ready Security Program - Aggressive Time Table (DRAFT)

For the CISO who needs to walk into a room Monday morning with a plan. This is meant as a quick reference to facilitate strategy and action. We assumed an aggressive time table in our recommendations, which may prove unrealistic for all organizations.

As you read through the table below, consider:

- Organization size, complexity, and budget should be taken into consideration. From very complicated environments to entirely SaaS-based ones, some would find it difficult to be agile, while others won't have available budget.
- Program updates should be considered in context, as some recommendations could prove contradictory if followed as-is. For example, the requirement to delay patching due to supply chain risks, with a cooldown period, directly competes with the need to patch faster. This calls for nuance in decision making and in policy, broadly, as well as in mitigating controls or specific incidents.

	Action	Category	Risk	Start	Horizon	What It Means
1	**Point Agents at Your Code and Pipelines**	Risk Control	CRITICAL ⚠⚠⚠	This week	Ongoing	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. • Commercial: Claude Code Security from Anthropic, Codex Security from OpenAI. • Open source: OpenAnt from Knostic, raptor (Claude Code framework), the exploitation-validator agentic skill, and agentic skills from Trail of Bits.
2	**Require AI Agent Adoption**	Operational Enabler	CRITICAL ⚠⚠⚠	This week	Ongoing	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.

3	**Defend Your Agents**	Risk Control	CRITICAL	This month	45 days	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.
4	**Establish Innovation, Acceleration Governance**	Governa-nce	CRITICAL	This week	6 months	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.
5	**Prepare for Continuous Patching**	Risk Control	CRITICAL	This week	45 days	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.
6	**Update Risk Models and Reporting**	Governa-nce	CRITICAL	This week	45 days	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.
7	**Inventory and Reduce Attack Surface**	Risk Control	HIGH	This month	90 days	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.

8	**Harden Your Environment**	Risk Control	HIGH	This month	6 months	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.
9	**Build a Deception Capability**	Risk Control	HIGH	Next 90 days	6 months	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.
10	**Build an Automated Response Capability**	Risk Control	HIGH	Next 90 days	12 months	Improve detection engineering and incident response capabilities to be systemic and, to the degree possible, autonomous. Examples: asset and user behavioral analysis, pre-authorized containment actions, and response playbooks that execute at machine speed.
11	**Stand Up VulnOps**	Risk Control	CRITICAL	Next 6 months	12 months	Long-term, there is no alternative to building a permanent Vulnerability Operations (VulnOps) function, staffed and automated like DevOps, but for autonomous vulnerability research and remediation. Owns continuous discovery of zero-day vulnerabilities across your entire software estate (from your own code to third-party software), and establishes automated remediation pipelines. Design VulnOps around triage discipline from the start.

ⓘ **Risk:** CRITICAL = immediate exposure if unaddressed; HIGH = significant exposure within 45 days; **Category:** Governance = structural prerequisite; Risk Control = direct risk reduction; Operational Enabler = makes risk controls executable

V. Executive and Board Briefing: the AI Risk Summary

Mythos is now a boardroom concern, and that creates an opportunity. This section is a working tool for CISOs preparing a leadership and/or board update, organized around two things: justifying the current program and making the case for what comes next. Every organization is different, so make sure you align the talking points and timelines with your actual current situation and programs.

↗ The Shift

AI at the capability level demonstrated by Mythos will transform how organizations operate, compressing development cycles and accelerating time to market. The business is already pursuing that value with current highly-capable models.

That same capability in adversary hands compresses the time between a vulnerability existing and causing business disruption from weeks to hours; a permanent acceleration, not a temporary spike.

This has two implications for the organization. First, several assumptions behind current risk metrics may no longer hold and need re-examination. We have moved into a world of containment and a focus on resilience, so metrics should now focus on the speed to recover to normal operations. Second, the same AI capabilities that create this risk also create a defensive opportunity: organizations can now identify their own weaknesses before attackers do, review code at machine speed, and respond to incidents faster than any human team can. Organizations that invest will be both faster to market and more resilient to attack.

Talking Point: AI Accelerates Both Sides

AI is making us faster and more competitive. But those same capabilities make attackers faster and more dangerous. It has compressed the time to a serious incident from weeks to hours, and that gap will continue to narrow. Turned inward, these tools let us find and fix our own weaknesses before adversaries do. Without attention to buying down risk, we move faster as a business while accumulating risk just as rapidly.

The security program this company has funded is what makes our AI security strategy viable. The investments already in place ensure that no single point of entry becomes a full business disruption. In an environment where entry points and weaknesses are discovered faster, that containment architecture is more valuable, not less.

With continued support, the changes we recommend here will return risk to pre-Mythos levels and demonstrate due diligence in response to a documented shift in the threat environment. This program builds the foundation that lets the business move fast with confidence.

Talking Point: An Aggressive Plan Is Needed

The funded foundation is why our program can adapt rather than react in a crisis. What has changed is the speed and volume it must absorb.

This is not an open-ended AI initiative. We are seeking alignment to execute a targeted and aggressive 90-day plan with clear owners and outcomes:

- **Increase People and Capacity.**
  Plan for repurposing of existing staff (within the security org, but also, and especially, within engineering teams) and/or onboarding of additional headcount and contractor capacity to handle the anticipated increases in triage, remediation, and incidents, while protecting experienced staff from burnout, especially as the first wave of Glasswing patches hits.

- **Deploy AI Tooling.**
  Formalize AI agent usage across all security functions as standard practice: scanning our own code, ensuring AI-driven review before code ships, and augmenting teams with purpose-built agents. This equips our teams to operate at the same speed as adversaries.

- **Harden Infrastructure.**
  Prioritize updating asset inventories; reducing unnecessary exposure; and enforcing segmentation, Zero Trust, egress filtering, and phishing-resistant authentication. Validate these elements across internal systems and key third-party providers (MSPs, SOCs).

- **Accelerate Procurement and Governance.**
  Align across functional teams (security, legal, engineering) to evaluate threats and fast-track priority defensive technology onboarding. Current approval cycles are too slow for the coming threat environment.

- **Update Playbooks.**
  Update technical and communications response plans to execute at the required speed and scale, including pre-authorized containment and coordination for simultaneous incidents.

- **Track Progress.**
  Provide regular check-ins throughout the 90-day period to capture results and identify roadblocks.

VI. Conclusions and Recommendations

AI-based attacks represent a structural shift in how offense and defense work, and it will not reverse. The cost and capability floor to exploit discovery is dropping, the time between disclosure and weaponization is compressing toward zero, and capabilities that previously required nation-state resources are now becoming broadly accessible.

While vulnerability discovery capabilities comparable to Mythos have shown to be present through earlier AI models, the Mythos announcement has grabbed the attention of the boardroom. Defenders can seize this opportunity and make a compelling business case to become “Mythos-ready” and prepare for an oncoming onslaught of patches.

Being “Mythos-ready” means:

- Engineering a resilient architecture that limits the ability of attackers to exploit discovered vulnerabilities and contains the impact if they are exploited.
- Discovering more vulnerabilities yourself in advance of any adversary (or vendor advisories).
- Responding quickly to incidents at scale and containing the impact to minimize business disruption.
- Accelerating your security program and staff capabilities with AI agents.

Empower your teams to use AI for defense, starting today. Every action in this brief can begin this week.

We have done this before. Y2K was a systemic threat with a hard deadline, and the industry met it through coordinated, disciplined effort. This is the same kind of problem, requiring the same kind of response, with more powerful tools available to defenders.

Building a “Mythos-ready” security program is not about reacting to one model or announcement. It is about permanently closing the gap between how fast vulnerabilities are found and how fast your organization can respond.

Appendix A: Historical Precedent

Background

This all began with the DARPA Cyber Grand Challenge, a landmark competition organized by DARPA in 2016 that demonstrated the potential of fully automated cybersecurity systems. Teams developed autonomous platforms capable of identifying, exploiting, and patching software vulnerabilities in real time, without human intervention. The challenge highlighted a shift toward machine-speed cyber defense, showing how automation and artificial intelligence could significantly enhance vulnerability management and incident response, while also raising important questions about trust, control, and the future role of human operators in cybersecurity.

By mid-2025, XBOW, an autonomous offensive security company, topped the HackerOne leaderboard. The DARPA AI Cyber Challenge (AIxCC) found 54 vulnerabilities in four hours of compute. Google's Big Sleep discovered real zero-days in open source.

Anthropic was used to automate full attack chains from reconnaissance through exfiltration. And, open source tools such as raptor proved autonomous vulnerability research is available to anyone able to use an agent.

In September 2025, Heather Adkins (CISO, Google) and Gadi Evron (CEO, Knostic) published a warning that attackers were racing toward a singularity moment, with autonomous vulnerability discovery and exploitation roughly six months away.

In February 2026 Anthropic, using Claude Opus 4.6, reported more than 500 high-severity vulnerabilities in open source software. AISLE found 12 OpenSSL zero-days, including a CVSS 9.8 vulnerability dating to 1998.

Linux kernel maintainers saw vulnerability reports climb from 2 to 10 per week, largely hallucinated at first, but that changed rapidly. The volume has held steady, but the reports are now all verified as real bugs.

The curl project, which originally discontinued its bug bounty program because it was drowning in hallucinated vulnerability reports ("AI slop"), last week echoed the above observation from the Linux team, reporting an increasing number of AI-supported quality security reports.

Sysdig documented an AI-based attack that reached admin-level access in eight minutes. This week, Gambit released a report on the AI-led compromise of Mexican government infrastructure, originally reported in February.

In March, Sergej Epp and others introduced the Zero Day Clock, visually demonstrating the disappearing time to exploit development, demonstrating the drastic fall in time to exploitation to less than a day in 2026. It is worth noting that the historical collapse in time-to-exploit has not yet produced a proportional increase in the impact of exploitation. Many of the most consequential incidents of recent years involved credential abuse, social engineering, or supply chain compromise rather than zero-day exploitation. The Zero Day Clock trend is a leading indicator of where attacker capability is heading, not a direct measure of current damage.

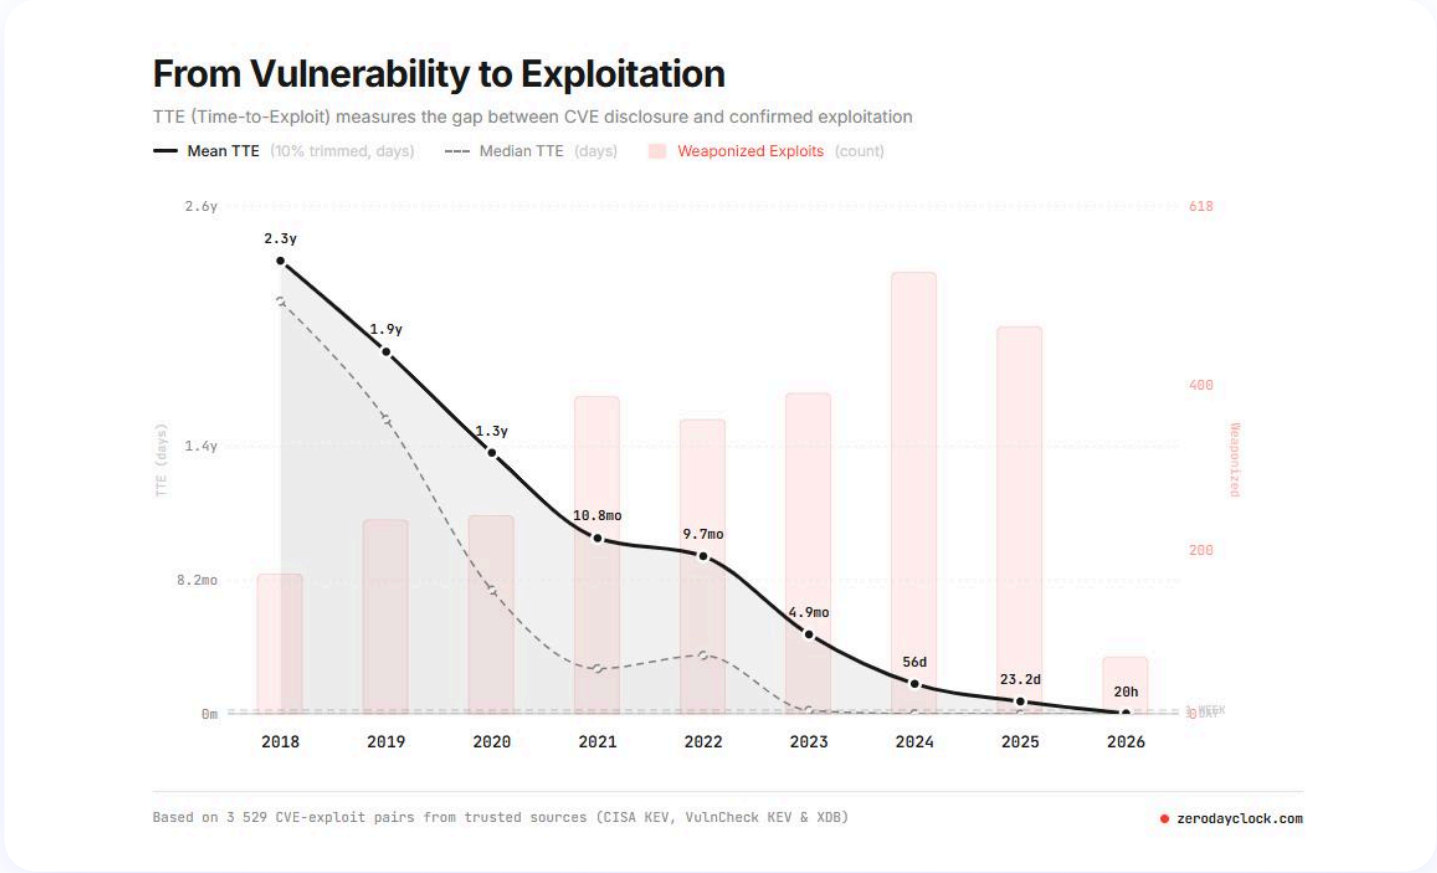

Diagram from the Zero Day Clock, by Sergej Epp, demonstrating the collapsing time to exploitation, which is now down to hours.

Appendix B: Mythos Risk Register Legend

OWASP LLM 2025 · OWASP Agentic 2026 · MITRE ATLAS · NIST CSF 2.0

1. Framework Prefixes

Every code in the Frameworks column belongs to one of these four frameworks.

LLMxx

OWASP Top 10 for LLM Applications 2025
Risks in LLMs used as application components

ASIxx

OWASP Top 10 for Agentic Applications 2026
Risks in autonomous AI systems that plan and act

AML.Txxxx

MITRE ATLAS
Adversarial techniques targeting AI/ML systems

GV.xx

NIST CSF 2.0 - Govern (GV)
Governance: context, risk strategy, roles, supply chain

ID.xx

NIST CSF 2.0 - Identify (ID)
Asset management, risk assessment, improvement

PR.xx

NIST CSF 2.0 - Protect (PR)
Access control, platform security, resilience

DE.xx

NIST CSF 2.0 - Detect (DE)
Continuous monitoring, adverse event analysis

RS.xx

NIST CSF 2.0 - Respond (RS)
Incident management and communication

2. All Framework Codes Used in This Register

Code	Full name and framework
AML.T0000	ML Model Reconnaissance - MITRE ATLAS
AML.T0018	Backdoor ML Model - MITRE ATLAS
AML.T0040	ML Inference API Access - MITRE ATLAS
AML.T0043	Craft Adversarial Data - MITRE ATLAS
AML.T0047	ML-Enabled Product Abuse - MITRE ATLAS
AML.T0051.000	LLM Prompt Injection (Direct) - MITRE ATLAS
AML.T0051.001	LLM Prompt Injection (Indirect) - MITRE ATLAS
ASI01	Agent Goal Hijack - OWASP Agentic Top 10 2026
ASI02	Tool Misuse and Exploitation - OWASP Agentic Top 10 2026
ASI03	Identity and Privilege Abuse - OWASP Agentic Top 10 2026

ASI04	Agentic Supply Chain Vulnerabilities - OWASP Agentic Top 10 2026
ASI06	Memory and Context Poisoning - OWASP Agentic Top 10 2026
ASI08	Cascading Failures - OWASP Agentic Top 10 2026
ASI10	Rogue Agents - OWASP Agentic Top 10 2026
LLM01	Prompt Injection - OWASP LLM Top 10 2025
LLM02	Sensitive Information Disclosure - OWASP LLM Top 10 2025
LLM05	Improper Output Handling - OWASP LLM Top 10 2025
LLM06	Excessive Agency - OWASP LLM Top 10 2025
LLM08	Vector and Embedding Weaknesses - OWASP LLM Top 10 2025
DE.AE	Adverse Event Analysis - NIST CSF 2.0 Detect
DE.CM	Continuous Monitoring - NIST CSF 2.0 Detect
GV.OC	Organizational Context - NIST CSF 2.0 Govern
GV.OV	Oversight - NIST CSF 2.0 Govern
GV.RM	Risk Management Strategy - NIST CSF 2.0 Govern
GV.RR	Roles, Responsibilities, and Authorities - NIST CSF 2.0 Govern
GV.SC	Supply Chain Risk Management - NIST CSF 2.0 Govern
ID.AM	Asset Management - NIST CSF 2.0 Identify
ID.IM	Improvement - NIST CSF 2.0 Identify
ID.RA	Risk Assessment - NIST CSF 2.0 Identify
PR.AA	Identity Management, Authentication, and Access Control - NIST CSF 2.0 Protect
PR.IR	Infrastructure Resilience - NIST CSF 2.0 Protect
PR.PS	Platform Security - NIST CSF 2.0 Protect
RS.CO	Incident Response Communication - NIST CSF 2.0 Respond
RS.MA	Incident Management - NIST CSF 2.0 Respond

3. Severity

Level	Meaning
Critical	Immediate exposure or increased risk if unaddressed
High	Significant exposure or increased risk within 45 days
Medium	Organizational risk requiring structured attention; does not create direct exploitable exposure but weakens the effectiveness of higher-priority controls if left unaddressed.

4. Risk Type

Type	Definition
Threat	External actor capability - controls raise cost but cannot eliminate it
Vulnerability	Internal exploitable condition - addressable through remediation
Capability gap	Defensive function missing or operating below the required level
Governance	Organizational or structural failure that amplifies every other risk