LTM

POV

# Claude Mythos and Project Glasswing

## What it means for your enterprise and the 90-day plan to get ready

**Yogesh Sharma**
Senior Director – Cybersecurity, LTM

# Table of Contents

**73%**
Expert CTF Tasks Mythos Solved (AISI Independent Test)

**90 Days**
Days Glasswing Public Report Timeline (From 8 April 2026)

**$100 Million**
Anthropic-Committed Glasswing Model Credits

# Overview

## Mythos is a capability inflection, not a product launch

Anthropic has kept Claude Mythos Preview behind closed doors precisely because it demonstrated autonomous discovery and exploitation of zero-day vulnerabilities in major operating systems and browsers. The UK AI Security Institute (AISI) confirmed that Anthropic's Claude Mythos Preview solved 73% of expert-level capture-the-flag (CTF) cybersecurity tasks, a feat unmatched by previous models.

## Project Glasswing puts a 90-day clock on the enterprise

Anthropic has committed to a public report within 90 days of the 8 April 2026 announcement, produced with AWS, Apple, Broadcom, Cisco, CrowdStrike, Google, JPMorgan Chase, the Linux Foundation, Microsoft, NVIDIA and Palo Alto Networks. The report will set industry guidance on vulnerability disclosure, patching, supply-chain and secure-by-design and will compress the window between disclosure and attacker weaponization.

## Enterprise response needs to have a focused plan

This POV aims to examine the impact of Claude Mythos Preview and Project Glasswing. Clients who treat the Mythos announcement as a targeted readiness exercise will be the ones who absorb the Glasswing report as a managed event rather than a crisis.

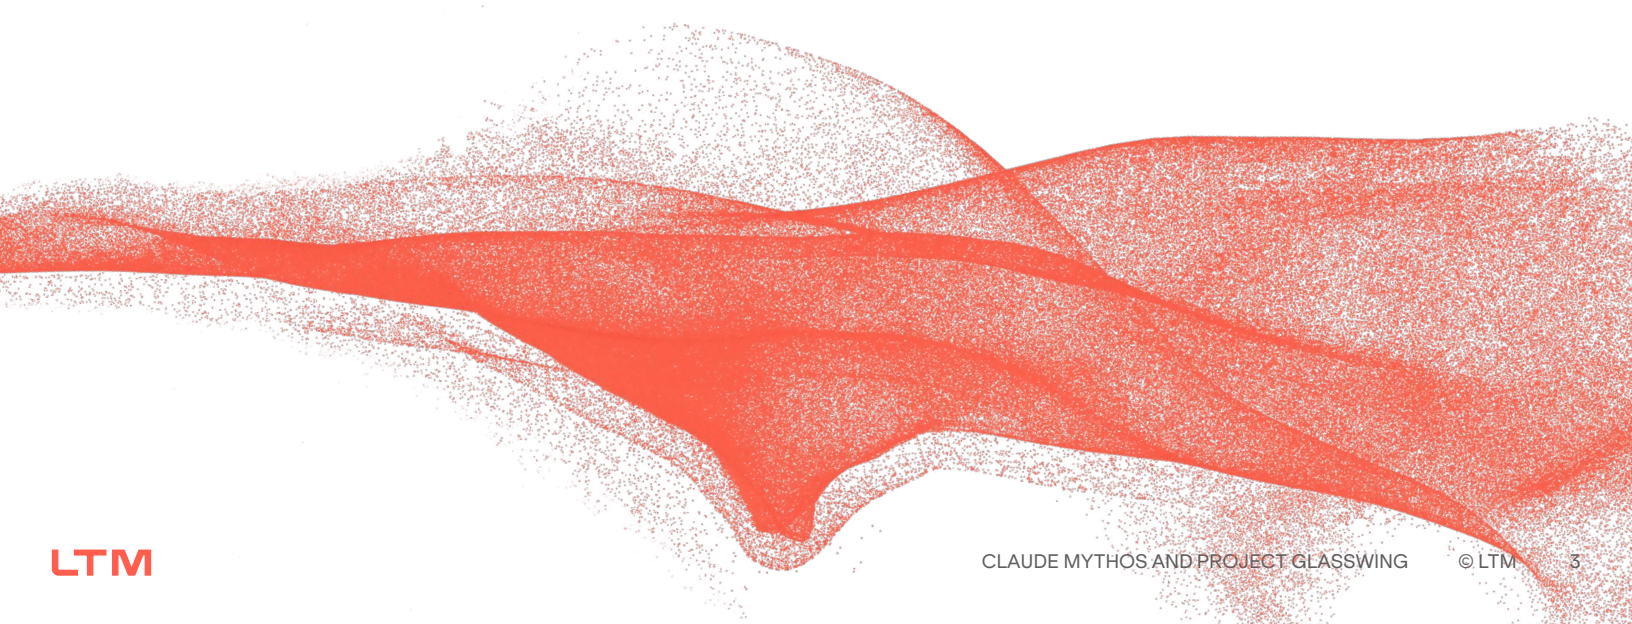

# What Mythos is—The Facts

Mythos is Anthropic's most powerful LLM to date capable of discovering vulnerabilities across major OS and browser platforms and autonomously writing working exploits in just a few hours. It is currently withheld from public release due to its offensive cyber capabilities.

Every claim below is drawn from Anthropic's announcement, the UK AI Security Institute's independent evaluation, and reporting by Outlook Business, Mashable, Forbes, CyberScoop and UKAuthority between 8 and 17 April 2026.

| Dimension | What Anthropic has Disclosed |
|---|---|
| **The model** | Claude Mythos Preview – Anthropic's frontier general-purpose LLM. Announced 8 April 2026. Not for public release. |
| **The claim** | Autonomous discovery of "thousands of high-severity vulnerabilities," including a 27-year-old OpenBSD bug, a 16-year-old FFmpeg flaw, and a chain of Linux kernel vulnerabilities enabling full system control. |
| **The test evidence** | In Anthropic's internal testing, Mythos produced working exploits 181 times out of several hundred attempts, versus near zero for Claude Opus 4.6. |
| **Independent validation** | UK AI Security Institute: 73% solve rate on expert Capture The Flag (CTF) tasks no prior LLM completed; averaged 24 of 32 steps on the "Last Ones" corporate-network simulation (prior ceiling: 16); first model to complete the full 32-step chain end-to-end. |
| The coalition | Project Glasswing – AWS, Apple, Broadcom, Cisco, CrowdStrike, Google, JPMorgan Chase, Linux Foundation, Microsoft, NVIDIA, Palo Alto Networks. USD 100 million Anthropic-committed credits plus open-source security donations. |
| The deliverable | Public report within 90 days covering disclosure, patching, supply-chain, secure-by-design and regulated-industry standards expected as early as July 2026. |
| Commercial path | Post-preview, Mythos-class access expected via Claude API, Amazon Bedrock, Google Cloud Vertex AI and Microsoft Foundry. No hyperscaler has published a date. |

# What Changes Because of Mythos

Exploitation activities that once took significant time, effort, and expertise can now happen much faster because of Mythos. This shortens the window organizations have to respond and shifts the focus of cyber risk toward speed of response and clear prioritization, rather than detection alone.

## A named Set of Open-source Components is Under Pressure

Mythos findings publicly cited so far in OpenBSD (27-year-old bug), FFmpeg (16-year-old issue), chain of Linux kernel vulnerabilities enabling full system control are not rare or specialized technologies. They ship inside mainstream enterprise platforms
and supplier products. Your exposure question is now asset and supplier-specific, not abstract.

## A 90-day Industry Clock Now Exists

Glasswing will publish industry guidance on disclosure, patching, supply-chain and secure-by-design within 90 days of 8 April 2026. Attackers will reverse-engineer the public findings; defenders will be expected to respond against that cadence. This is the single most consequential enterprise timeline to plan against in 2026

## Standards Will Be Shaped By a Specific Coalition

The Glasswing participants -AWS, Apple, Broadcom, Cisco, CrowdStrike, Google, JPMorgan Chase, Linux Foundation, Microsoft, NVIDIA, Palo Alto Networks -will set the first reference point for how regulated industries are expected to handle Mythos-class disclosures. Influence over the downstream standard sits with them, not outside observers.

## The Announcement Has Already Triggered a Supervisory Response

Within days, the US Federal Reserve Chair and Treasury Secretary convened banking CEOs on Mythos. Expect supervisors in other regulated sectors to ask comparable questions inside the 90-day window.

# What It Means – Industry Viewpoint

Claude Mythos changes the scale and speed of cyber risk, but its business impact differs sharply by industry. The following table highlights how this shift will materialize across key sectors.

| Sector | Where Mythos-class Capability Changes Your Exposure |
|---|---|
| **Banking and Capital Markets** | The Federal Reserve System and Treasury emergency meeting with banking CEOs is a direct signal. Expect supervisors to ask Tier-1 banks to stress-test core banking, trading and clearing estates against Mythos-style disclosure clusters. JPMorgan Chase's Glasswing seat means standards will be shaped from inside BFSI. |
| **Insurance and Healthcare** | Mythos specifically found deep flaws in long-lived open-source components. EHRs, medical devices and insurance platforms that ship FFmpeg, OpenBSD-derived stacks or kernel modules will see disclosures land inside regulator-notified windows under HIPAA, the EU European Health Data Space and India's DPDP. |
| **Energy and Utilities (OT/ ICS)** | AISI's OT cooling-tower scenario failed -but in the IT segment feeding OT, not in OT itself. That is not a reassurance: Mythos demonstrably accelerates the IT-to-OT crossing path that attackers already exploit. Emergency patching in OT is physically gated, so disclosure-to-remediation windows become safety events. |
| **Manufacturing and Hi-Tech** | FFmpeg and embedded Linux kernels -both named in Mythos findings -propagate through OEM firmware on quarter-long rebase cycles. A Glasswing disclosure wave lands during those cycles, not after them. |
| **Retail and CPG** | Glasswing's public report is expected in early July 2026 -inside many retail patch-freeze windows that run through back-to-school and festival planning. The disclosure cadence, not the vulnerability count, is the commercial risk. |
| **Telecommunications** | Carrier and service-provider estates sit on the same open-source components Mythos already tested. Autonomous lateral-movement capability specifically compresses the time state-level actors need once they are inside. |

# Our Five Convictions on Mythos

These five convictions reflect where Mythos challenges conventional thinking and where organizations must shift focus to stay ahead. They are less about new tools, and more about changing how we prioritize, operate, and scale security.

## Detection is Not the Binding Constraint

The real challenge isn't visibility; it's prioritization and throughput. With hundreds of valid findings, the question becomes: what gets fixed first, and can the patch pipeline keep pace?

## Clean up theLegacy Estate we've been Ignoring

Take action on systems older than 10 years – modernize, isolate or retire. The 27-year-old and 17-year-old bugs Mythos found weren't flukes – they're the norm in old code.

## Build an AI-grade Security Operations Center

Move from signature-based detection to behavior-based, AI-assisted monitoring. Invest in 24x7 threat hunting – assume breach, hunt continuously, don't wait for alerts.

## Fix the Software Supply Chain

Most Mythos-cited components reach you through third parties. Mandate a Software Bill of Materials (SBOM) for every application.

## Apply AI Deliberately and Securely to Scale Operations

Ensure robust security and governance by protecting data, enforcing compliance, and regularly auditing AI systems to minimize risks and build trust.

# Prioritizing and Fixing the Flood – LTM's Playbook

When the Glasswing report publishes, every enterprise of meaningful size will face the same problem at the same time: a surge of valid, high-severity findings mapped to components they own, with patches of varying maturity and a finite change-management bandwidth. This section provides best practices that answer "which first, and how fast?" in a form you can operate against.

## A Four-factor Prioritization Score

Rank each incoming Mythos and Glasswing finding on four factors.

| Sector | Where Mythos-class Capability Changes Your Exposure | How it Drives Action |
|---|---|---|
| **Reachability** | Is the vulnerable code path actually reachable in our deployment and configuration? | Unreachable = deprioritize immediately. VEX "not affected" statements belong here — use them. |
| **Exploitability** | Is there a working or likely public exploit? Is it on CISA KEV? . | Known exploit + reachable = top of queue. Common Vulnerability Scoring System (CVSS) alone is insufficient; exploitability context is decisive. |
| **Blast radius** | Is this a crown-jewel system, a regulated workload, or a lateral-movement pivot point? | Same Common Vulnerabilities and Exposure (CVE) can be Tier 0 on a core banking node and Tier 2 on an internal wiki. |
| **Compensating controls** | Is a Web Application Firewall (WAF) rule, segmentation boundary, Endpoint Detection and Response (EDR) detection or IdP control already in place? | Good compensation buys time — it does not remove the item from the register. |

A finding that is reachable and exploitable on a crown-jewel asset with no compensating control is Tier 0 regardless of CVSS. A high CVSS finding on an unreachable code path with a compensating WAF rule is Tier 3.

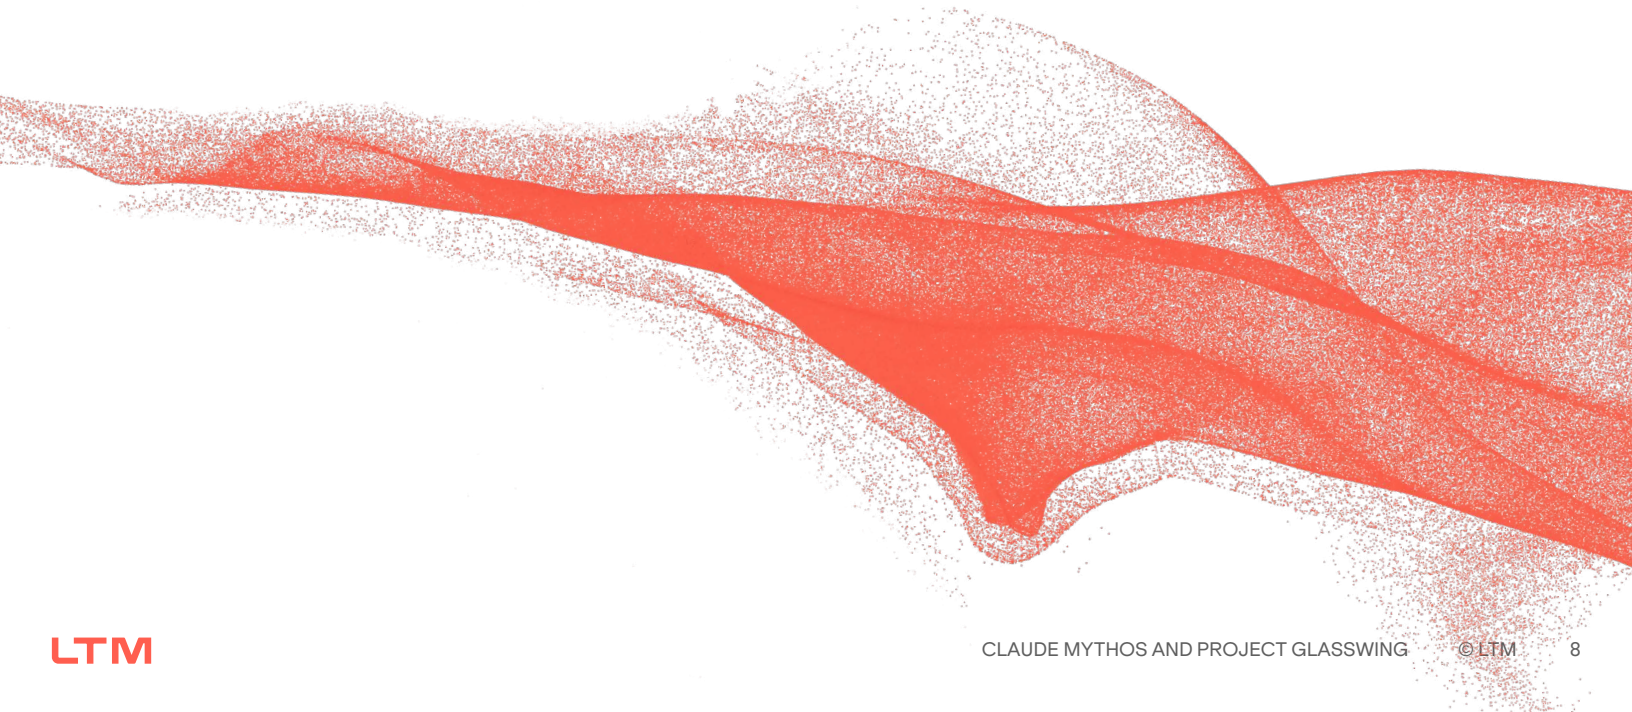

## A Tiered Remediation SLA

Pre-agree the SLAs with the audit committee now, not after the report drops. Public commitment to these tiers is what unlocks emergency change windows.

| Tier | Profile | Target SLA | Governance |
|---|---|---|---|
| **Tier 0** | Reachable, exploitable, crown-jewel, no compensating control | Patch or mitigate within 72 hours | CISO-owned; emergency change window; daily stand-up until closed |
| **Tier 1** | Reachable and exploitable, non-crown-jewel | Patch within 14 days | VulnOps-owned; standard expedited change |
| **Tier 2** | Reachable, no known exploit | Patch within 30 days | Service-owner responsibility; scheduled change |
| **Tier 3** | Unreachable or fully compensated | Track in exception register; review at 60 days | Named owner, compensating control, expiry date |

## Remediation-capacity Plays

- Auto-patch lane for low-blast-radius code. Canary + automated rollback. Use it for stateless services, internal tooling and non-customer-facing batch — the majority of your estate by count, a small share by risk.
- SBOM plus VEX to filter noise. Machine-readable VEX ("Vulnerability Exploitability Exchange") flags let suppliers and internal teams publish "not affected" assertions against specific CVEs removing them from the queue without losing audit trail.
- Pre-approved disclosure-wave change windows. Secure standing CAB approval now for emergency patch windows triggered by the Glasswing publication event. Debating windows in the moment is how enterprises miss the first 72 hours.
- Exception register with expiry dates. Every deferred patch moves to a register with named owner, compensating control, and mandatory re-review date. The register is a first-class artefact at the audit committee.
- Shift-left supplier attestation. Before your team touches a CVE, require the supplier's VEX statement or patched build. Inbound work is reduced by removing vendor-resolvable items from your queue entirely.
- Reserve 20% of remediation capacity. Hold back one in five patching cycles from BAU for the 90-day window. Organizations that absorb disclosure waves are the ones that refused to run their patch pipeline at 100% in normal times.

# A 90-day Operating Plan (Glasswing-Aligned)

This operating plan is designed so that Week 10–12, the expected arrival of the actual Glasswing report, is a structured step, not a scramble.

| Horizon | Focus | Specific Outcomes to Demand |
|---|---|---|
| **Weeks 1–2** | Mythos exposure mapping | Inventory systems that ship the component families Mythos specifically targeted -OpenBSD-derived stacks, FFmpeg, vulnerable Linux kernel versions -across owned estate and top suppliers. Board / audit-committee readout. |
| **Weeks 3–4** | Glasswing intake cell | Inventory systems that ship the component families Mythos specifically targeted -OpenBSD-derived stacks, FFmpeg, vulnerable Linux kernel versions -across owned estate and top suppliers. Board / audit-committee readout. |
| **Weeks 4–6** | Supplier attestations | Formal requests to your top-50 suppliers for SBOM refresh and written attestation on Mythos-cited component families. Escalate non-responders to procurement and legal. |
| **Weeks 5–7** | Patch-velocity rehears | Rehearse an emergency patch cycle on a representative OpenBSD / FFmpeg / Linux kernel workload end-to-end: identify, qualify, canary, roll out, rollback. Capture mean-time-to-patch; this becomes the Glasswing KPI. |
| **Weeks 6–8** | Detection retune | Update SOC detection content for Mythos-style lateral-movement patterns observed in the AISI 32-step simulation. Validate coverage with an AI-assisted purple-team exercise. |
| **Weeks 7–9** | Regulator posture | Prepare holding statements and regulator-facing responses aligned to the jurisdictions that apply to you (DORA / NIS2 for EU; Fed / OCC for US BFSI; CERT-In six-hour directive and DPDP notification windows for India). |
| **Weeks 8–10** | 72-hour Glasswing drill | Tabletop a Friday-evening publication of the Glasswing report containing CVEs that hit three of your top-20 systems. Measure triage time, patch decision latency, external-comms readiness. |
| **Weeks 10–12** | Report and institutionalize | Absorb the actual Glasswing report. Board readout on findings, actions taken, residual risk and lessons-learned. Move the intake cell from temporary to standing. |

# Where to Be Skeptical

## The "Thousands of Zero-days" Figure is Vendor-reported

Anthropic has not disclosed false-positive rates, a comparison to existing static-analysis and fuzzing tools, or how much human triage sits behind the number (Heidy Khlaaf, AI safety engineer, in Mashable, 14 April 2026). Treat it as directional, not audited.

## Commercial Motive is a Real Confound

The overlap between the safety narrative and Anthropic's fundraising cycle is genuine. Plan against the capability direction, not the press-release integers.

## Hyperscaler Distribution is Unconfirmed

No hyperscaler has published a date for Mythos-class availability on Bedrock, Vertex AI or Foundry. Do not let procurement plans bake in a quarter.

## OT impact is Ambiguous

AISI's OT cooling-tower test failed in the IT segment. That is neither proof of safety nor of catastrophe but it means OT owners should still assume Mythos accelerates the IT-to-OT path.

# What We Are Watching Next

The developments below will determine how fast Mythos moves from headline to day to day business impact.

- Whether the Glasswing report includes CVSS-level disclosures or stays at the standards / process layer.
- Whether any Glasswing partner confirms a Mythos-class availability date on Bedrock, Vertex AI or Foundry.
- Whether cyber insurers convert Mythos-related exposure to explicit exclusion or affirmative coverage with governance preconditions.
- Whether other frontier labs publicly demonstrate comparable capability, which would change the containment equation.

# Conclusion

Claude Mythos and Project Glasswing mark a structural shift in cybersecurity, rather than a transient technology cycle. The demonstrated ability of an AI system to autonomously discover and exploit complex vulnerabilities compresses the traditional boundaries between discovery, weaponization and impact. What was once sequential is now concurrent and faster. For enterprises, the implication is not simply more vulnerabilities, but a different operating environment. The combination of a known set of exposed component families, a defined 90-day disclosure horizon, and a coalition that will shape industry standards creates a rare moment of predictability within disruption. The decision is clear: treat Mythos as a contained, time-bound preparedness exercise or face it later as an unmanaged event. Organizations that prepare in advance will manage the disclosure wave with control; those that do not will encounter it under pressure, with greater operational and regulatory risk.

# About the Author


## Yogesh Sharma

Senior Director – Cybersecurity, LTM


Yogesh is a Practice Head for application security, vulnerability management, and AI Security CoE, leading enterprise security strategy and execution across modern applications, cloud-native platforms, and AI-driven systems. He brings deep expertise in scaling DevSecOps, advancing risk-based vulnerability management, and embedding security into the engineering lifecycle, enabling organizations to achieve resilient, secure, and accelerated digital transformation at enterprise scale.

# References


1. Assessing Claude Mythos Preview's cybersecurity capabilities – 7 April 2026: https://red.anthropic.com/2026/mythos-preview/
2. Outlook Business -"Anthropic's Claude Mythos Breakthrough: Project Glasswing Launched to Prevent AI Cyber-Crisis" -8 April 2026.: https://www.outlookbusiness.com/deeptech/anthropics-claude-mythos-breakthrough-project-glasswing-launched-to-prevent-ai-cyber-crisis
3. Mashable -"Is Anthropic's Claude Mythos a big stunt, or a real security threat?" -14 April 2026: https://in.mashable.com/tech/108482/is-anthropics-claude-mythos-a-big-stunt-or-a-real-security-threat-what-the-experts-say
4. CyberScoop -"Here's how cyber heavyweights in the US and UK are dealing with Claude Mythos" -13 April 2026: https://cyberscoop.com/claude-mythos-ai-cybersecurity-threat-report/
5. UKAuthority -"Claude Mythos Preview demonstrates advanced cybersecurity capabilities" -17 April 2026: https://www.ukauthority.com/articles/claude-mythos-preview-demonstrates-advanced-cybersecurity-capabilities
6. Forbes (Tim Keary) -"Claude Mythos Isn't Just A PR Stunt, It's The New Face Of Offensive AI" -15 April 2026: https://www.forbes.com/sites/timkeary/2026/04/15/claude-mythos-isnt-just-a-pr-stunt-its-the-new-face-of-offensive-ai/
7. UK AI Security Institute -independent evaluation of Claude Mythos Preview, April 2026: https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities