BLACK DUCK®

EXECUTIVE BRIEF

YOUR MYTHOS MANDATE: MODERNIZE APPSEC FOR MACHINE-SPEED EXPLOITS

FRONTIER AI MODELS ARE COLLAPSING THE TIME FROM VULNERABILITY DISCOVERY TO EXPLOITATION. THE PRIORITY IS NO LONGER DETECTION—IT'S COVERAGE, VISIBILITY, AND RESPONSE AT SPEED.

Patrick Carey, Vice President of AppSec Solutions

When Anthropic first announced Claude Mythos and Project Glasswing in April, the overwhelming response by security teams was, understandably, "Let's wait and see." The headlines were alarming, the stakes were high, and the threats were vague. And to some extent, we've seen this before—when Generative AI tools first entered the enterprise, when Log4Shell exposed gaps in software visibility, and when supply chain mandates forced structural change. In each of these cases, initial response wasn't immediate action but cautious evaluation.

But Mythos is genuinely different, and it changes the playbook that those past decisions were based on. Speed is now the defining variable in application security—and for most organizations, it's the point of failure. Mythos didn't just improve vulnerability discovery; it collapsed the time between discovery and exploitability. What once took weeks or months now happens in hours. The headlines focus on the *scale* of what Mythos can find. The real impact is *how quickly* those findings can be turned into working attacks.

This shift exposes a fundamental weakness in modern AppSec programs. Most of these programs were built for a world where time was an asset. Teams could assess, triage, and remediate within a meaningful window. That is no longer the case. When exploit timelines move at machine speed, any gap between finding and fixing presents risk.

The mandate is clear: Close the gap. Organizations that can match the speed of discovery with the speed of response will reduce risk. Those that can't are operating at a structural disadvantage that compounds with every new vulnerability identified.

WHY MYTHOS IS DIFFERENT

AI models have been getting better at finding vulnerabilities for a while now. That's not new. What makes Mythos different is its combination of three capabilities that together represent a step change in what an attackers can do, and how little skill they need to do it.

- **Mythos generates working exploits without scaffolding.** In internal Anthropic lab tests, Mythos produced 181 functional exploits of Firefox. The previous model, Claude Opus 4.6, produced just two under the same exact conditions. Now, exploit development is scalable and repeatable, and timelines are compressed beyond what most AppSec programs are designed to handle.
- **Mythos can identify and link complex vulnerabilities.** It can combine memory corruption, logic flaws (such as auth bypass and privilege escalation), and cryptographic weaknesses into multistep exploit chains. It can also analyze and reverse-engineer closed-source binaries, enabling end-to-end vulnerability discovery and chaining that previously required elite, well-resourced attackers.
- **Mythos does significantly more with a single prompt.** It doesn't need elaborate configuration or agent setup. This gives "one-shot capability" to an attacker, which reduces the time required to produce a working exploit from weeks to mere minutes.

Since Anthropic announced Mythos Preview on April 7, alongside Project Glasswing, the model has already identified thousands of zero-day vulnerabilities across every major operating system and browser, not just Firefox. Glasswing, which gave critical infrastructure providers, industry partners, and open source maintainers early access to patch their own software before general availability, surfaced more than 10,000 high- or critical-severity security flaws just from the initial cohort of 50 or so partners.

Anthropic took the responsible approach by controlling the Mythos rollout. It worked with the Linux Foundation, Microsoft, Cisco, and other major software producers to give them a head start. But within 6 to 12 months, many other AI companies are likely to have Mythos-class models—and they could release them without any of those safeguards. What happens then will be a very different situation.

WHAT LLM-GRADE ATTACKS LOOK LIKE

It's easy to tune out the abstract claims of what these models can do. But when you see what the attacks look like in practice, the implications become stark and the mandate clear. Here are four recent data points to consider.

- **The Firefox number.** Producing 181 working exploits compared to just two by the previous version, under identical lab conditions, should give you immediate pause. The implication for your environment is that an attacker using a Mythos-class model can generate and test exploit candidates against a target codebase at a volume and speed that eliminates any meaningful response window. If your current patching cycle takes weeks, the math doesn't work.
- **The Sysdig compromise.** In February 2026, Sysdig documented a real-world AI-based attack that reached admin-level access in eight minutes. Think about your mean time to detect. Is it under eight minutes? For most organizations, the answer is no—and the delta is where the damage happens.

- **The Zero Day Clock.** In March 2026, Sysdig CISO Sergej Epp launched the Zero Day Clock to visualize the change in time-to-exploit (TTE) over the past several years. The data is based on 3,529 CVE-exploit pairs from CISA KEV, VulnCheck KEV, and XDB. In 2018, mean TTE was 2.6 years. Now, it's under one day. The traditional patch cycle was built for a world where you had months between disclosure and weaponization. It's woefully inadequate in a world where attackers can find vulnerabilities and exploit them at machine speed.
- **Democratization of the attacker pool.** The broader issue is bigger than Mythos itself. It's what happens when these capabilities spread. The DARPA AI Cyber Challenge in June 2025 found 54 vulnerabilities in four hours of compute time across 54 million lines of code. XBOW became the first autonomous system to top HackerOne's U.S. leaderboard, outperforming all human hackers on the platform. Google's Big Sleep discovered 20 real-world zero-days in open source projects, and all of them were found autonomously. Open source tools like GitHub Raptor already enable autonomous vulnerability research for anyone who can run an off-the-shelf agent. All these examples point to the same conclusion: Capabilities that previously required the resources of a nation state are now broadly accessible. The asymmetry between attackers and defenders—attackers only need one win and defenders need continuous perfection—is getting worse, not better.

THE NVD PROBLEM COULDN'T HAVE COME AT A WORSE TIME

There's another piece to this that hasn't gotten enough attention. At exactly the moment your teams need the most reliable, timely vulnerability intelligence to respond to the Mythos-driven surge in disclosures, the primary public infrastructure for delivering that intelligence changed the rules.

In April, NIST announced a major operational change to the National Vulnerability Database (NVD). CVE disclosure volumes have essentially tripled over the last five years, and NIST doesn't have enough staff to keep up. So it decided to fully enrich only CVEs that align with its obligations to the U.S. government going forward. Everything outside that scope will still be published, but it will be tagged "Lowest priority—not scheduled for immediate enrichment."

In practice, this means the NVD will fully enrich only the vulnerabilities that make it into the Known Exploited Vulnerabilities (KEV) subregistry, which accounts for less than 10 percent of everything reported. The other 90 percent will be handed off to certificate naming authorities (CNAs). Some CNAs will do excellent work. But many won't—which means the fidelity of that enrichment will be wildly uneven.

The timing is brutal. You're about to face a flood of vulnerability disclosures from the Glasswing shockwave. And then the other AI model vendors will start their own discovery programs. At the same time, the public data source most security programs depend on for prioritization guidance is stepping back from the vast majority of disclosures.

This creates a two-sided problem: The number of vulnerabilities requiring attention is going sharply up, while the quality and timeliness of the data you need to prioritize and remediate them is going sharply down. For organizations that don't adapt, it's not a manageable imbalance—it's a structural failure waiting to happen.

WHAT YOU NEED TO DO TO BE MYTHOS-READY

It's clear that immediate action is required. But being "Mythos-ready" isn't about reacting to one model or announcement. It's about permanently closing the gap between how fast you find vulnerabilities and how fast you can respond to them. Here's how to break that down into actionable priorities.

Fill in the Gaps and Double-Down on Your AST Program

You can't protect software you aren't scanning and managing. That has always been true, but in a Mythos environment, uncovered applications and repositories aren't just a gap in your reporting, they're the first place an AI-powered attacker is going to look.

Many organizations ship software that hasn't been fully tested. As AI-generated code makes it easier for teams outside of traditional development departments to spin up new projects, that problem compounds fast. Your first priority is getting everything tracked and tested—uniformly, not on a team-by-team basis where individual groups decide what level of scrutiny they want to apply to their own software.

Automation is the only way this is possible at scale. Your security policies need to apply consistently across every branch, every pull request, and every build, without requiring someone to remember to run a scan. You should also update your prioritization framework. In a world of AI-driven exploits, a lower-severity vulnerability with a confirmed exploit in the wild is a higher priority than a theoretical critical vulnerability with no observed exploitation.

Manage Your Supply Chain Like Your Life Depends on It

When new vulnerabilities are exposed (and we're about to see a lot of them), bad actors immediately start looking for organizations that haven't patched the affected components. Your Software Bill of Materials (SBOM) is only useful if it's complete and current.

False negatives in supply chain security almost always come from the same source: incomplete SBOMs. You miss a vulnerability because you don't know the component is there. Black Duck® SCA was built to close this gap, covering modern languages, compiled languages like C and C++, binaries, AI-generated code, and AI models pulled from Hugging Face and similar repositories.

There's also a growing threat that traditional SCA wasn't designed to catch: malicious components such as lookalike packages, compromised committer behavior, and AI-generated malicious code injected into open source projects. Because they're zero-days by design, these don't have CVEs. Your supply chain security program needs active monitoring capabilities that go beyond CVE-based scanning—it must be able to detect suspicious package behavior and anomalous activity before a vulnerability is even disclosed.

Start Moving Toward a Secure Agentic SDLC—but Don't Abandon the Deterministic Foundation

As you evaluate your defensive strategy, you'll undoubtedly ask yourself: *If models like Mythos are so good at finding vulnerabilities, why not just use an LLM to secure my code?*

It's a fair question. But it's based on a misunderstanding of what Mythos actually is, and what it takes to get the results you need.

The Anthropic findings that generated all the headlines—the Firefox exploits, the OpenBSD analysis—were produced by highly trained security researchers doing a lot of configuration, context-setting, and intense analysis. The OpenBSD research project found dozens of findings, true. It also cost $20,000 in tokens. That's not continuous security testing. And it's not a model you can drop into a CI/CD pipeline and expect to run for every commit across every repository you have.

Mythos is a model, not a platform. It's not designed for the kind of continuous, automated, policy-driven security testing that an enterprise AppSec program requires. LLM-based analysis is probabilistic—results vary between runs, hallucinations happen, and there's no auditability that satisfies a regulator or enterprise compliance team.

What moving to an agentic SDLC actually means is using AI to automate the find-prioritize-fix workflow. It means taking the output of your scanning, applying smart triage, generating fix recommendations and pull requests, and validating the results at machine speed. AI is an accelerant for your existing capabilities, not a replacement for them.

Deterministic SAST and SCA remain the foundation. They deliver consistent, reproducible, auditable results—the kind that hold up in compliance reviews. They track every dependency to its license obligation and supplier risk in a way that LLMs can't replicate. And they run continuously at a cost structure that scales.

The organizations that understand this will be the ones that build a hybrid architecture: deterministic scanning as the auditable base, with LLM-powered capabilities layers on top for deeper analysis and automated remediation.

Stop Relying Solely on the NVD for Vulnerability Intelligence

If your vulnerability prioritization and remediation workflow is still primarily based on NVD enrichment, you have a dependency problem that you need to address now. For roughly 90 percent of disclosed vulnerabilities, the enrichment you've relied on won't show up in time—not in a world where exploits can be weaponized in hours.

Black Duck Security Advisories (BDSAs) are how we address this for our customers. The Black Duck Cybersecurity Research Center (CyRC) team monitors more than 2,500 security feeds and data resources, maps findings across over 10 million tracked open source components, and publishes complete BDSAs nearly three weeks ahead of the corresponding NVD entry on average. During active zero-day events, BDSAs are reviewed and updated hourly, a cadence the NVD never achieved and is now moving further away from.

Beyond timing advantages, BDSAs provide what NVD enrichment has historically lacked: temporal scoring. BDSAs account for exploit availability and remediation status, provide contextual tags like "Zero-click RCE" and "Malicious code identified" for instant triage, deliver project-aware remediation guidance that points you to the specific fixed version and viable workaround for your specific environment, and ensure coverage for more than 3,500 vulnerabilities you won't find in the NVD at all.

HOW BLACK DUCK HELPS YOU BECOME MYTHOS-READY

In addition to BDSAs, Black Duck offers a host of solutions to help you prepare for Mythos and other LLM-grade security attacks.

Black Duck Polaris™ Platform closes the AST coverage gap at enterprise scale. It synchronizes automatically with your source code repositories—GitHub, GitLab, Bitbucket, Azure DevOps—so nothing new goes undetected. Event-based test automation ensures that every pull request and build triggers the right security assessment. Security policy is set centrally and applied uniformly. And AI-powered remediation assistance starts closing the gap between finding and fixing.

Black Duck SAST and SCA solutions provide the complete supply chain defense and deterministic scanning foundation your program requires. Our SCA tools build the complete SBOMs you need to catch every vulnerable component, including those in AI models pulled from repositories like Hugging Face. The Black Duck KnowledgeBase™ tracks 250,000 unique vulnerabilities across 10 million open source components, and Black Duck Bridge enables auto-fix pull requests, eliminating the manual triage loop between findings and remediation. Black Duck SAST tools provide the deterministic, auditable static analysis needed for compliance and enterprise-scale reliability, which no LLM-based tool can replicate.

Black Duck Signal™ is an agentic application security layer that enables you to make the shift from diagnostic AppSec to active defense. Signal integrates LLM-enhanced code analysis, agentic test-triage-fix workflows, AI-powered threat modeling, active supply chain defense, and agentic penetration testing—all powered by Black Duck's proprietary ground truth, ContextAI™. Its active supply chain defense goes beyond CVE-based scanning to detect malicious lookalike components, suspicious committer behavior, and anomalous supply chain activity that traditional SCA isn't designed to catch. And the proprietary data we've accumulated through decades of SCA, SAST, and DAST gives Signal a contextual intelligence that no competitor can replicate.

START NOW: THE WINDOW IS CLOSING

Mythos is neither a cause for panic nor an excuse to freeze. It's a clear indicator that the assumptions underlying most AppSec programs—how much time you have, what public vulnerability data is available, and what skill level an attacker needs—have fundamentally changed. The organizations that treat this as an opportunity to modernize their programs and are willing to engage, close the gaps, and lean into the capabilities available to them are going to come out ahead.

We built Black Duck's portfolio—comprehensive, deterministic AppSec combined with AI-powered acceleration—for exactly this moment.

Talk to a Black Duck specialist today about becoming Mythos-ready

About Black Duck

Black Duck® meets the board-level risks of modern software with True Scale Application Security, ensuring uncompromised trust in software for the regulated, AI-powered world. Only Black Duck solutions free organizations from tradeoffs between speed, accuracy, and compliance at scale while eliminating security, regulatory, and licensing risks. Whether in the cloud or on premises, Black Duck is the only choice for securing mission-critical software everywhere code happens. With Black Duck, security leaders can make smarter decisions and unleash business innovation with confidence. Learn more at www.blackduck.com.