seemplicity

/solution brief

Closing the Find-to-Fix Gap at the Speed of Claude Mythos

Claude Mythos has fundamentally broken traditional security by moving the threat landscape from human-speed discovery to autonomous, agentic-speed exploitation. As the window between discovery and weaponization collapses from weeks to minutes, your primary risk is no longer the finding itself but the organizational lag in responding to it.

Seemplicity is the only Agentic Exposure Action Platform built to counter machine-speed threats with machine-speed remediation. By replacing manual coordination with an automated, agentic defense, we ensure that your team outpaces AI-driven attacks by turning every discovery into an immediate, verified fix.

The Mythos Threat

Traditional vulnerability management processes were built for a world where the gap between discovery and exploitation was measured in weeks or months. Models like Claude Mythos have the potential to compress this window into minutes.

The primary risks introduced by Mythos-class AI include:

- **Timeline Compression:**
  Models like Mythos collapse the time between vulnerability discovery and exploit development, making manual coordination an obsolete strategy.
- **The Signal Deluge:**
  AI-driven discovery creates a massive spike in findings. The challenge is to validate and fix it before an autonomous agent can exploit it.
- **Lowered Barrier to Entry:**
  Engineers without formal security training have used AI to generate working exploits overnight, expanding the pool of sophisticated threats exponentially.
- **The "Unfixable Gap":**
  Even when a security tool identifies a risk at AI speed, the organizational drag of manual ticket routing and unclear ownership keeps the window of exposure open for days or weeks.

Traditional Vulnerability Management Tools Are Now Obsolete

Legacy vulnerability management was built on a single, comforting assumption: Exploitation is hard. For years, security teams prioritized only the top 10% of vulnerabilities, confident that the manual effort required for exploit development would keep the remaining 90% "safe" in the backlog. In the age of AI, that assumption is no longer true. Risk-Based Vulnerability Management (RBVM) was built on assumptions that are no longer valid in the age of AI-driven exploitation. When AI can weaponize CVE descriptions in seconds, relying on historical data and manual complexity assessments isn't just insufficient; it's a liability.

The following table highlights the critical disconnect between legacy security assumptions and the reality of AI-driven adversaries.

Legacy Vulnerability Management Assumption	The Mythos Reality
No Public PoC: Logic suggests no public Proof of Concept equals low likelihood of exploit, allowing for deferred patching.	**Automated Exploit Generation:** Mythos generates working exploits directly from CVE descriptions. A public PoC is no longer a prerequisite for an attack.
High Attack Complexity: Vulnerabilities with high CVSS complexity are often deprioritized as "too hard" for most attackers.	**Infinite Skill at Zero Cost:** AI removes the barrier of human effort. Mythos possesses infinite technical skill at zero marginal cost; complexity is no longer a defense.
KEV Listing: Security teams use the CISA KEV (Known Exploited Vulnerabilities) list as a primary driver for urgency.	**Lagging Indicators:** KEV is reactive. By the time a vulnerability is officially listed as "actively exploited," AI-driven attackers have already moved on to the next target.
EPSS Scores: Low EPSS scores (e.g., < 5%) are frequently relegated to the backlog as "unlikely to be exploited."	**Historical Bias:** EPSS is trained on historical human behavior. It cannot predict the speed or scale of AI-assisted exploit generation.

Defeating Agentic Threats with Agentic Defense

Seemplicity is the **Agentic Exposure Action Platform** designed to close the gap between finding and fixing. We provide the agentic defense necessary to counter agentic threats.

1. **Bridge the "Find-to-Fix" Gap at Machine Speed**
   While Claude Mythos finds vulnerabilities in hours, Seemplicity ensures your remediation teams act in minutes. By integrating with your existing scanners (like Wiz, CrowdStrike, and Tenable), Seemplicity automatically ingests findings, enriches them with business context, and routes them to the correct owners in Jira or ServiceNow. This eliminates the manual handoffs that AI-driven attacks exploit.

2. **Agentic Prioritization for a Deluge of Data**
   You cannot fix everything at once, especially when AI is generating thousands of new findings. Seemplicity's AI agents proactively analyze business risk to prioritize your backlog. We consolidate scattered, individual findings into consolidated, actionable fixes, allowing your team to focus on the exposures that matter most to your specific environment.
3. **Investigation for Root Cause Resolution**
   While others stop at the finding, Seemplicity performs an investigation. Our agents perform the heavy lifting of tracing code reachability and mapping the blast radius. Instead of simply forwarding a raw CVE number and a deadline, we provide the exact root cause and the asset-specific guidance required to resolve the issue.
4. **Evidence-Based Verification**
   In a world of machine-speed exploits, simply closing a ticket is a liability when AI can re-exploit an improperly patched asset in seconds. Seemplicity uses independent, automated verification to confirm the window of exposure is truly shut. We do not trust the ticket status; we trust the evidence.
5. **Operational Proof at Scale**
   Security leaders need to move beyond reporting to prove resilience against autonomous threats. Every decision, priority change, and exception in our platform is backed by an automated reasoning chain. This creates a full audit trail that allows you to prove to your board and auditors that exposure is being managed in real time, not just cataloged.

Measurable Impact in the AI Era

Organizations using Seemplicity to manage their exposure see immediate, definitive results:

60%	80%	75%	98%
Reduction in Mean Time to Remediation (MTTR).	**Decrease** in manual operations.	**Increase** in SLA compliance.	**Backlog Reduction** via automatic deduplication and aggregation.

Security teams do not struggle with detection; they struggle with action. In a world where Claude Mythos is finding your vulnerabilities, Seemplicity ensures you fix them first.

Stop the discovery-fix gap today. Visit seemplicity.ai to learn more.

Trusted by the World's Leading Enterprises

Carlsberg Group · MORNINGSTAR · SoFi · Graphic Packaging International · NTT · WEST BEND · Hertz · OpenWeb