Anthropic Mythos & Why CYE

Sales CheatSheet – April 2026

What You Need to Know

- **Anthropic's Claude Mythos Preview was launched on April 7, 2026.** Mythos is an AI model that autonomously discovers and chains vulnerabilities at machine speed.
- **Project Glasswing** is a $100M restricted partner program using Mythos to identify systematic risk in critical software before attackers do. Partners, incl. Crowdstrike, Palo Alto Networks, Google, Microsoft, AWS.
- **Pro**: Claude Mythos has already proved it can discover 1000's of vulnerabilities (15+ years of age).
- **Con**: Mythos' greatest value is its greatest weakness – enabling attackers to exploit faster and more efficiently (as was reported by Bloomberg on Apr. 23)

Bottom line – amplifies the need for remediation: smart, in context auto-remediation.

Mythos: What it can and can't do

What it CAN do

- It is already uncovering **thousands of high-severity vulnerabilities** and complex attack paths.
- It can **outperform humans** at vulnerability discovery and chaining
- Discover and chain **multistep attack paths**, not just single exploits.

What it CAN'T do

- Map vulnerabilities to business context or prioritization tied to financial impact.
- Give a unified exposure view across the environment
- Offer defensible decisions that hold up at board level
- Enable/operationalise remediation.

The Risks

For Security Teams

Attackers can now use AI to find and exploit vulnerabilities at scale. The time between "unknown vulnerability" and "real-world attack" is collapsing.

- **False confidence** - incomplete view of true exposure
- **Misplaced focus** - chasing noise, missing business-critical risks
- **Unworkable remediation** - plans that can't be executed

For Cye and Other Vendors

- Agentic **AI could be argued as a replacement** for assessments, red teaming, and attack path analysis
- **Customers may question the need** for external security validation

Why Customers Need Cye More Than Ever

AI replaces vulnerability discovery - not prioritization, remediation or validation.
Reducing exposure with measurable outcomes is the real challenge for enterprises.
This is Cye's sweet spot.

Built on Data Nobody Else Has	**AI You Can Trust**
• Trained on **hundreds of thousands of real incident claims** and validated attack paths • Powers proprietary models built over a decade & impossible to replicate • Cye **i**dentifies which vulnerabilities, in which sequence, actually threaten the business	• Built to protect **customer safety and business continuity** – not a compliance checkbox • Combines assessment with continuous validation • Trust is demonstrated with evidence - not assumed
Cyber-Trained AI over Generic AI	**Defensible Decisions, Not Subjective Scores**
• **We don't just run AI - we continuously train it** on real-world mitigation outcomes • **Clear gap vs. generic AI** - not high-level advice, but environment-specific actions • **Execution-ready remediation** - tools, configs, dependencies, policy, governance	• Translates exposure into **financial business impact** • Enables **decisions that stand up to board scrutiny** • Provides **continuous proof that risk is being reduced**

AI creates the volume. Cye brings the precision, context, and remediation.

What to Say to Prospects & Customers

See What AI-Driven Attackers See

AI chains vulnerabilities into real attack paths, but those paths are generic and lack business context and most tools fail to show this. Cye shows which paths actually reach critical business assets – based on your specific environment.

Ask "Which 2-3 weaknesses, if chained together, could disrupt your business this quarter?"

Speak the Language of the Board

Traditional risk models are breaking in the AI era. Boards now face liability and governance issues if they fail to adopt AI-assisted defensive practices. Cye translates technical exposure into financial risk, enabling boards to accurately evaluate cyber decisions.

Say: "We show what reduces financial risk fastest – and how to prove it."

From Volume to Decisions

AI floods teams with issues and alerts - everything looks urgent, but signal is low Teams struggle to prioritize what matters and act fast enough. Cye cuts through the noise - prioritizing business-critical risks and driving clear, actionable decisions.

Say: **"*Not everything urgent is important — we show what matters and what to do next."***

Supply Chain & AI Agent Risk

AI agents are turning supply chain risk from a static dependency problem into an actively executed one — a single upstream weakness can now propagate across thousands of systems in minutes. Cye maps how these weaknesses translate into real attack paths, quantifies the business impact, and prioritises what to fix first.

Say: "It's not where the vulnerability is — it's where it leads."

Objection Handling

Q. "Does AI make Cye obsolete?"

A. No. Mythos proves discovery is no longer scarce. Cye exists to manage the consequences of that reality.

Q. "Why do we need Cye if we deploy agentic AI assessments."

A. Agentic AI assessments will highlight vulnerabilities. Cye highlights actionable remediation, continuously validated that you can trust.

Q. "Why do I need another platform?"

A. AI increases volume and speed. Without prioritization and financial context, teams drown.

Anthropic Mythos proves that cyber risk now moves at machine speed — Cye is how enterprises turn that reality into continuous, defensible exposure reduction.