CONTRAST SECURITY

SOLUTION BRIEF

Get Mythos ready: Secure your apps against AI

Stop cheap, automated attacks.

Anthropic's Claude Mythos Preview is an AI that autonomously finds and exploits software vulnerabilities. It found zero days hiding for decades, including a 17-year-old remote code execution flaw in FreeBSD. It built working Linux kernel privilege-escalation exploits for under $2,000 each. Every AI lab is building a version of this.

The threat is not theoretical. Unauthorized actors accessed Mythos on the day it was announced. Automated exploit development has removed the skill barrier that once slowed attackers down. Your existing tools were not designed for a world where that barrier no longer exists. Here is what to do about it.

$2k
Cost to build a working Linux kernel privilege escalation exploit with AI

73%
Mythos success rate on expert-level hacking challenge[1]

33%
Of all intrusions begin with exploiting a software vulnerability[2]

158 days
Average time for organizations to detect a breach[3]

Three steps. In this order.

01 OBSERVE

Know what's exploitable right now.

Mythos finds vulnerabilities that attackers can exploit. Scanning tools can't tell you which of your vulnerabilities are actually reachable and exploitable — runtime can.

Runtime instrumentation sees your application from the inside: which code paths run in production, which vulnerabilities are exposed and which are being probed.

Without that visibility, your team is triaging a CVE list with no idea what's in front of an attacker.

02 PROTECT

Block exploits before the patch ships.

Scanning finds issues. It does not block exploits. There is a gap between "we know about this" and "we are safe from this." That gap is where breaches happen.

Runtime protection closes it. Contrast blocks attacks at the point of execution, inside the running app, with no patch required and no CVE needed.

17-year-old unpatched vulnerability? Blocked. Zero day that was never cataloged? Blocked. Unknown until Mythos found it last night? Still blocked.

03 FIX

Make remediation strategic, not reactive.

Run your AI scan. It will find real issues. The problem is that it cannot tell you: which of those findings are exploitable, which are being actively probed and which can safely wait.

Runtime data answers those questions. You get a prioritized list, not a noise pile. AI-assisted fixes work on confirmed, exploitable issues — not static guesses.

That is where Contrast fits: not instead of scanning, but after it, giving it the context it cannot generate on its own.

AI scanning is a good start. Runtime is what makes it complete.

AI scanning tools are getting better. Run them. They will surface real vulnerabilities. But every scanning tool — AI or otherwise — shares one limitation: it analyzes code at a point in time, from the outside.

Runtime protection works differently. It instruments your application directly, living inside the running code, not outside it. It sees attacks as they happen and blocks them before data moves.

Here is what no scanner can tell you:

- Which vulnerabilities are actually reachable by an attacker right now, given how your application is deployed and running?
- Which are being actively probed? Mythos-class tools are already being used to test systems in the wild. Runtime sees that. Scanners do not.
- What happens when an attacker finds it before you patch it? A scan result does not block an exploit. Runtime does.

Scanning tells you what might be broken. Runtime tells you what is under attack right now — and stops it. That is the gap. That is where Contrast lives.

If someone is pitching you a pure AI scanning tool

The question to ask any AI scanning vendor	Why it matters
What's the all-in cost — not the API bill, but the triage and remediation labor for every finding your team has to chase?	AI-generated findings are cheap to produce and expensive to triage. More findings do not mean less risk.
If you run the same scan twice, what percentage of findings will match?	In Contrast Labs testing: 17%. You can't build a remediation strategy on findings that don't reproduce.
What happens when an attacker exploits a vulnerability before you patch it?	Their answer tells you everything. Contrast blocks that attack. They don't.

Note: Established SAST vendors with AI-assisted scanning are a different category. They bring years of rule development and language coverage that purpose-built AI scanners don't have. These questions are for tools that are AI-scanning only.needed to systematically test and deploy fixes on a standard schedule.

"Finding flaws is a solved problem. Defense is not, because defense requires context that lives nowhere in your codebase. Someone has to actually defend the thing."

— Dave Lindner, CISO, Contrast Security

"The real breakthrough is not AI alone. It is AI guided by runtime truth about how software actually behaves in production. That is why Contrast is so well positioned for where the market is going."

— Jeff Williams, Founder & CTO, Contrast Security

Try Contrast

[1] UK AI Security Institute
[2] Mandiant M-Trends 2025
[3] IBM Cost of a Data Breach Report 2025

Contrast Security is the world's leader in Runtime Application Security, embedding code analysis and attack prevention directly into software. Contrast's patented security instrumentation enables powerful Application Security Testing and Application Detection and Response, allowing developers, AppSec teams and SecOps teams to better protect and defend their applications against the ever-evolving threat landscape.

contrastsecurity.com

6800 Koll Center Parkway
Ste 235
Pleasanton, CA 94566
Phone: 888.371.1333