

# The CSA Mythos Response: Deploy Deception in 90 Days

AT-A-GLANCE

250+ CISOs have signed off on 11 priority actions—including building a deception capability

## What is Mythos?

Mythos is Anthropic's AI model capable of autonomously discovering thousands of critical vulnerabilities across major operating systems and browsers, generating working exploits without human guidance. In internal lab testing, Mythos generated 181 working exploits on Firefox where Claude Opus 4.6 succeeded only twice. The industry views Mythos as significantly increasing attacker capabilities, with the cost and skill floor for vulnerability discovery dropping significantly.

## The CSA's Response

Five days after Mythos was disclosed, the Cloud Security Alliance published an [emergency strategy briefing](#), with contributions from top security leaders including the former CISA Director, NSA Cybersecurity Director, and CISOs from companies including Google, Netflix, Wells Fargo and more. The briefing flagged “Inadequate Incident Detection and Response Velocity” as a CRITICAL capability gap due to the risk of “detection and response at human speed against machine-speed attacks.”

## Why Deception Is Mandated Among the 11 Priority Actions

CSA Priority Action #9 calls for deploying capabilities where detection is based on attacker behavior rather than the use of a known attack tool or CVE. This shift in methodology is a structural difference because signatures fail against novel exploits by definition—the vulnerability doesn't exist in a detection database. Behavioral detection fails when attackers use legitimate tools such as PowerShell, standard APIs, valid credentials—which are indistinguishable from normal employee activity.

[Deception does not rely on either](#), so if a decoy reports an interaction, that alert is necessarily confirmed—no investigation or triage needed, no false positives. Because a decoy alert carries zero ambiguity, deception is the only tool where automating containment is operationally sound—isolating hosts, revoking credentials, and blocking IPs at the speed of a Mythos-level, AI-driven attack.

### The “AI Vulnerability Storm”: Building a “Mythos-ready” Security Program

Expedited Strategy Briefing  
By the CSA CISO Community, SANS, [un]prompted,  
the OWASP Gen AI Security Project, and the wider community.

Contact [cloud@cloudsecurityalliance.org](#) with any inquiries.

Original release: 12 April, 2026  
Last updated: 01 May, 2026

Version 1.0

The latest version of this document can be found [here](#).

cloud security alliance SANS [un]prompted OWASP GenAI SECURITY

The Cloud Security Alliance convened 250+ CISOs and security leaders—including former CISA, NSA and White House cyber leaders—to help companies prepare for Mythos-enabled attacks.

**The result:** an emergency strategy briefing with **11 priority actions** for a Mythos-ready security program.

## #9: Build a Deception Capability

Classified as HIGH risk with a mandate to implement in the next 90 days.

*“Deception is attack-tool and vulnerability independent, identifying attacks and attackers based on their TTPs.”*