

If We Knew Then What We Know Now:

Influence of Mythos and Agentic AI on Assurance

Troy Leach

Executive Director, CSAI

Chief Strategy Officer, CSA

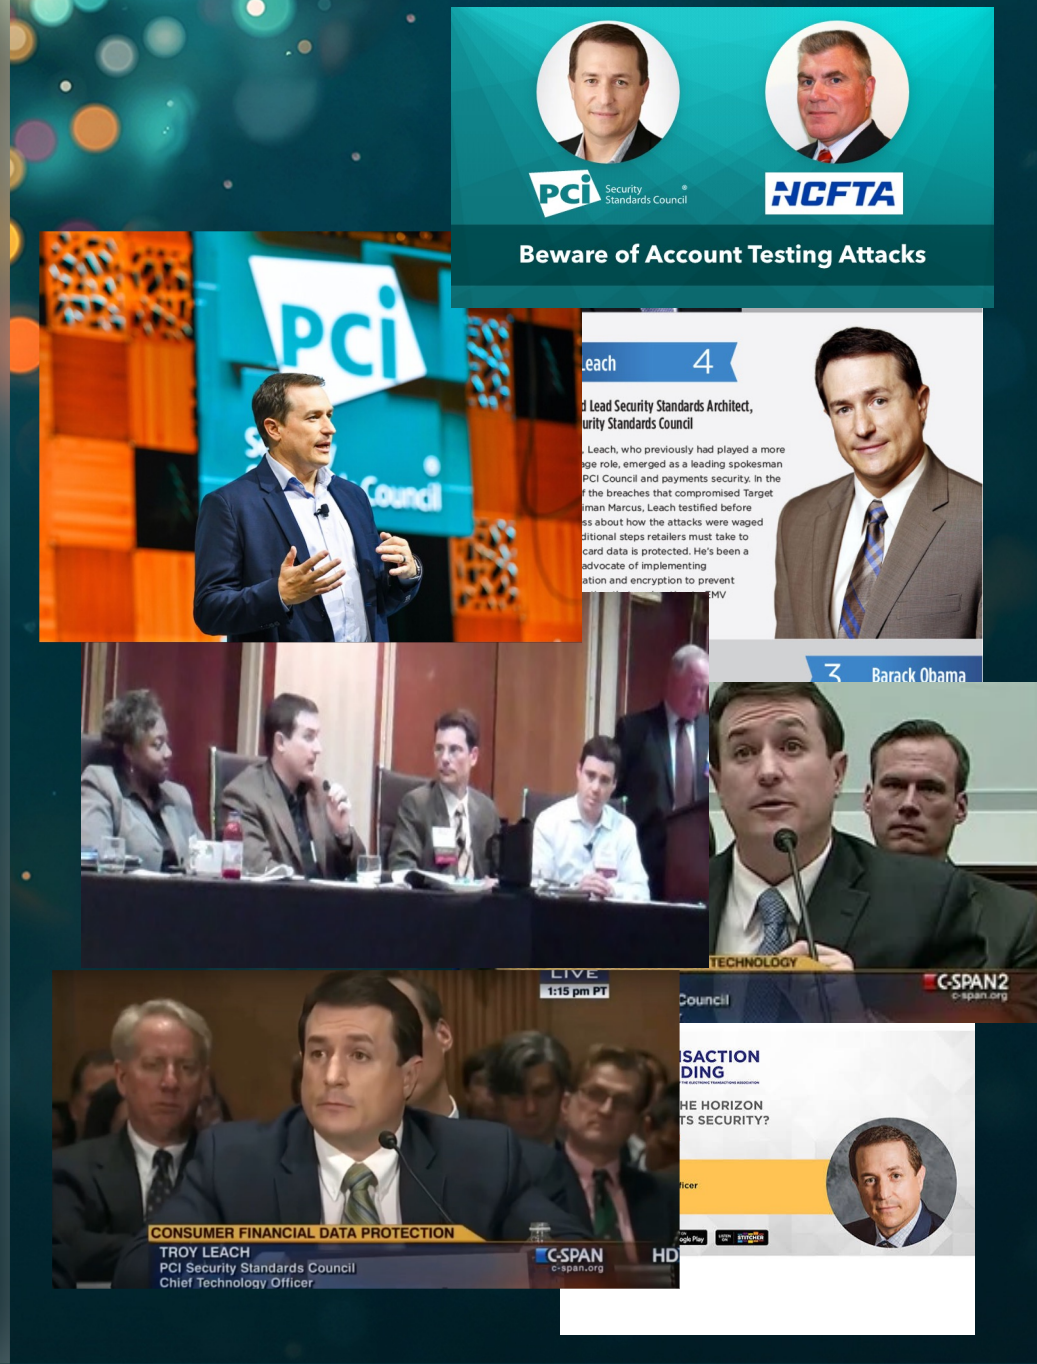


BACKGROUND

2004-2007: Helped to establish PCI Council and first TWG Chair

2008-2022: Chief Standards Architect and CTO for PCI SSC

Have collaborated on NIST, ISO, ANSI, CSA and other frameworks



Three Forces Reshaping Assurance

Cloud-native & ephemeral services

Infrastructure exists for minutes

AI Model Capabilities

Speed and scale to weaponize exploits

Autonomous Actions

Limited oversight and too much permission



Infrastructure is ephemeral

Transient workloads and ephemeral instances

Exists for minutes defeating point-in-time inspection.

Containers living under 5 minutes climbed from 20% (2019) to 70% (2024)

PCI DSS v4.0 rightfully emphasizes continuous compliance



Cloud Control Matrix v4.1 (CCM) mapped to PCI DSS v4.0

Mythos Was A Turning point

Exploits without scaffolding

Complex chained vulnerabilities

“One-shot” (single prompt) capability



OpenClaw and Hugging Face

Agentic AI Autonomy Challenges

Unsupervised agents

Supply-chain integrity

Agents are the company they keep



Intent Remains the Same

The threat model has evolved however

PCI DSS was intended to be outcome-based: understand the security objective, then validate

AI’s scale and cloud velocity is highlighting the importance of outcome-based intent



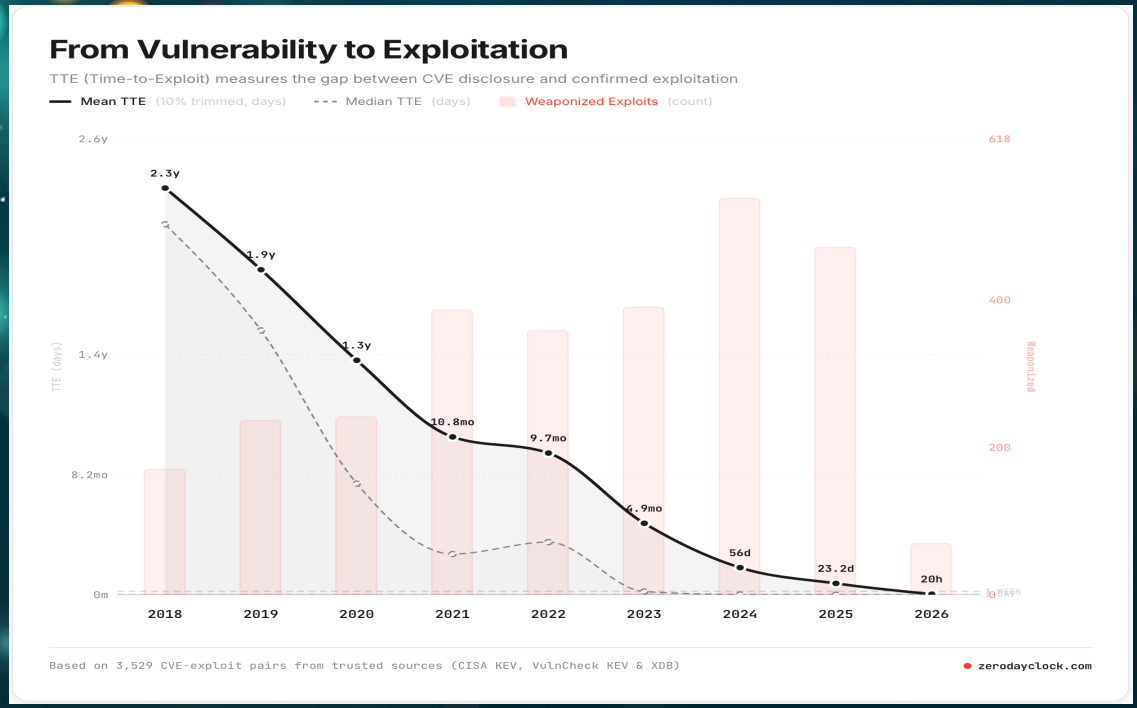
Things to Consider

Exploitation of discovered vulns has changed

Patch velocity must track the exploit clock, not the calendar

Agents defy segmentation (e.g. Hugging Face)

At 8 hours, identity hygiene — not patch speed — is the decisive control.



From 2.3 years in 2018 to 8 hours in 2026

Things to Consider

Behavioral analytics provides stability for AI

Then: Expensive, noisy and marginal risk reduction

Now: Autonomy, drift, alternative code push require new approach to visibility



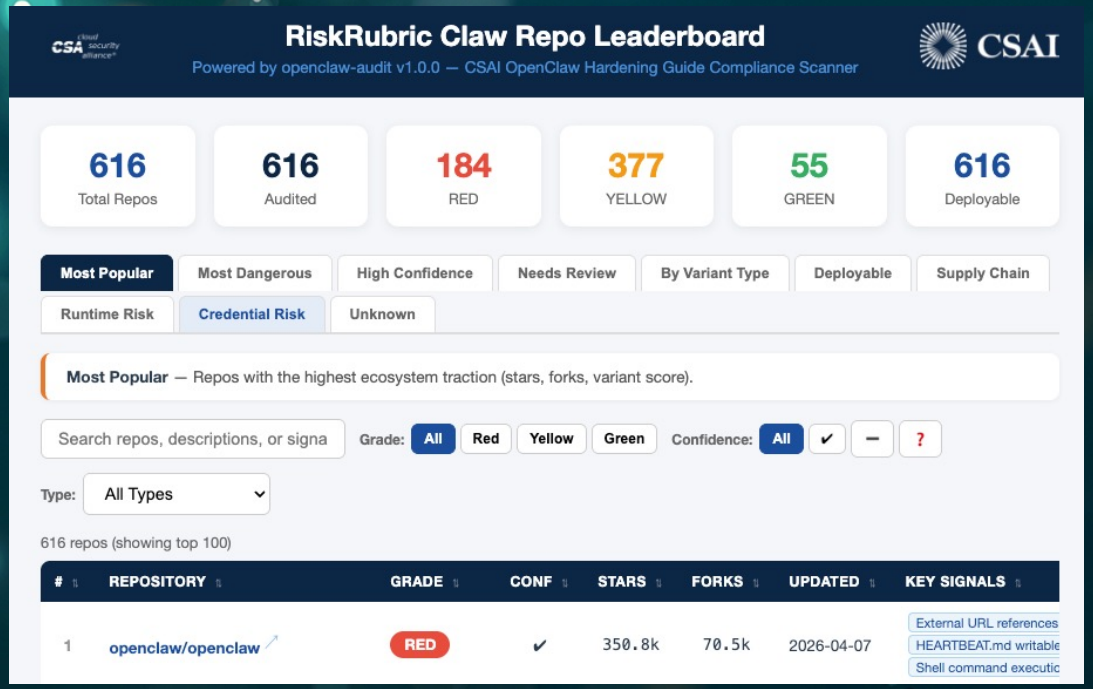
Third-party Risk

Managing the shifting dependency

Telemetry-driven, continuous monitoring

ML Classification for Policy Enforcement

Embedded AI in SaaS solutions



Risk Rubric – RiskRubric.ai

Things to Consider

Continuous Assurance

When infrastructure lives for minutes, the snapshot validates a system that's already gone.

Automated evidence collection can reduce audit effort by 60–80% (Splunk)

Continuous monitoring reduced audit findings by 50-70% (Sirion / Ponemon Institute)



Things to Consider

Non-Human Identities or Application/System Accounts

NHI outnumber (human) users 82:1

50% report NHI-related compromises (*ESG 2024*)

2% of permissions being used in container environment (*Sysdig*)

85% expect agentic commerce to be mainstream soon (*CSA State of Cloud and AI in Financial Services*)



Agentic Trust Framework

Grant agents autonomy = employees trust: earned, scoped, and revocable

Other Free Resources for Consideration

Example of tools found at labs.cloudsecurityalliance.org

AARM Autonomous Action Runtime Management

Runtime Accountability (aarm.dev)

MCP Server Top 10 Security Risks

Communication Protocol



Final Thoughts

Next generation of security and assurance will require change

The principles endure: know your data, control access, prove it works.

Observability, automation and retesting at the speed of your asset changes

Patch to the exploit clock, not the calendar.

Grant agents autonomy the way you grant employees trust

Whatever AI gives the attacker, institutionalize defensively first

Contact Information



Troy Leach

Executive Director, CSAI

tleach@csai.foundation

https://www.linkedin.com/in/troyleach/