Agentic AppSec Unleashed '26
by Checkmarx

The Rise of Agentic AI: Rethinking Security Programs

Caroline Wong
Chief Strategy Officer, Axari,
Author of AI Cyber Security Handbook

Two Near-Misses

Each time, we braced for catastrophe. Each time, it didn’t land the way we feared.

	WHAT WE FEARED	WHAT ACTUALLY HAPPENED
CVE / MITRE April 2025	Funding expires and the program goes dark — no new CVE IDs, national vulnerability databases deteriorate, and the tooling, threat intel, and incident response built on them start to break.	CISA extended the contract overnight; the catalog never went dark. A CVE Foundation formed, and a more durable funding arrangement followed — “no funding cliff.” The single-sponsor fragility, though, is now undeniable.
Mythos April 2026	Anthropic’s announcement: a model autonomously finding and exploiting zero-days at near-expert level — thousands of high-severity bugs, a multi-step network attack end-to-end. An “AI vulnerability storm” headed for every attacker.	Anthropic withheld public release — gated to ~40 orgs under Project Glasswing. And it’s not one superweapon: capability is jagged and distributed, with far smaller, cheaper models replicating much of it. The floor keeps rising.

Each time, the catastrophe didn’t arrive on schedule — so we exhaled. And the pattern kept compounding.

“

Mythos is just the new normal.

Daniel Miessler · April 2026

We Are Asking the Wrong Question

Same technology. Different reactions.

Overhyped

- We’ve seen this before
- Nothing fundamentally new
- Just another hype cycle

Transformational

- Everything just changed
- This changes the risk
- Harder for defenders

Operational

- How do we deploy this?
- How do we govern it?
- How do we scale it safely?

The technology doesn’t care what stage you’re in.

Two Curves That Collide

The risk is not the technology. It’s the gap.

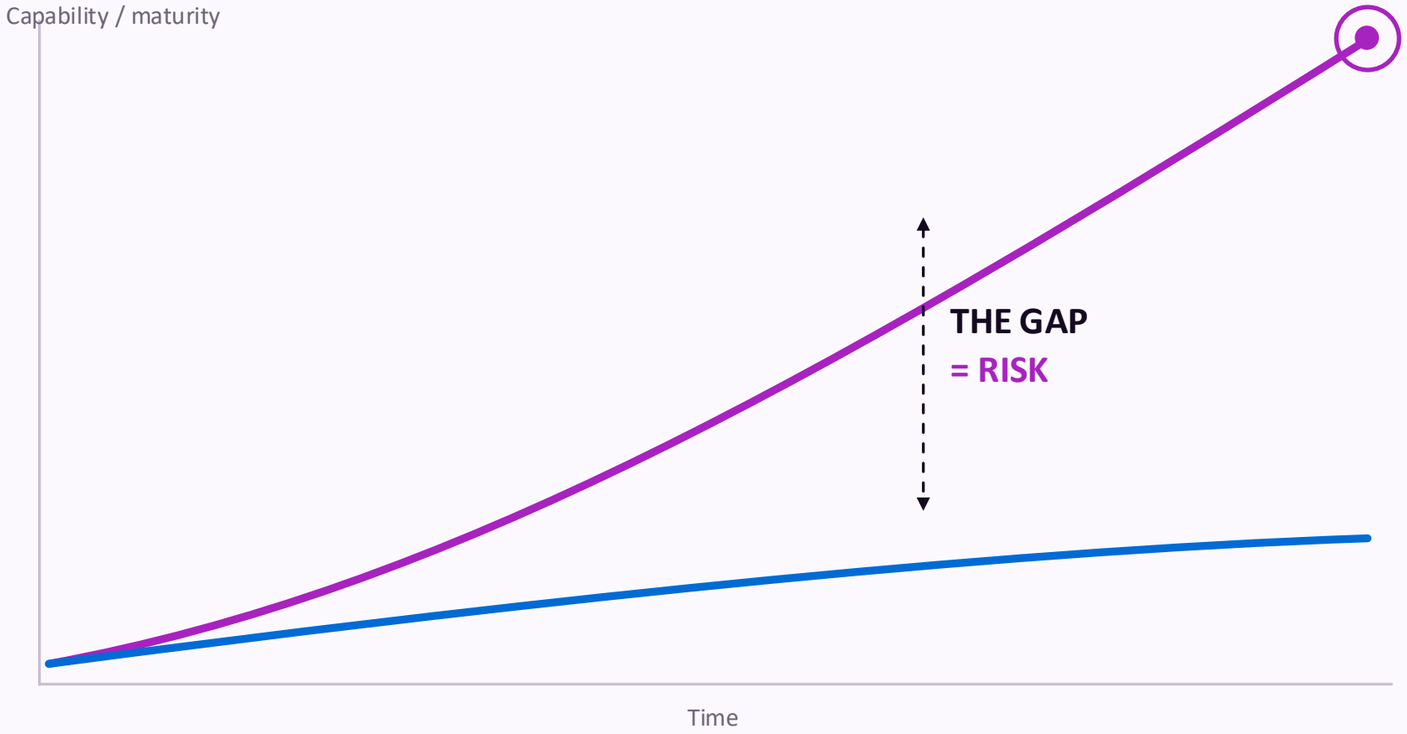

AI CAPABILITY

Accelerating exponentially

- Faster discovery
- Faster exploitation
- Automation & scale
- Autonomy at scale

ORGANIZATIONAL REALITY

Improving, but far more slowly

- Human capacity
- Process maturity
- Governance

AI is accelerating. | Organizations are not. | Risk lives in the gap.

AI vulnerability storm.

CSA CISO Community · SANS · OWASP AI Security Project · April 2026

AI cybersecurity after Mythos: the jagged frontier.

Stanislav Fort, AISLE · April 2026

The Collision

Risk lives in the gap.

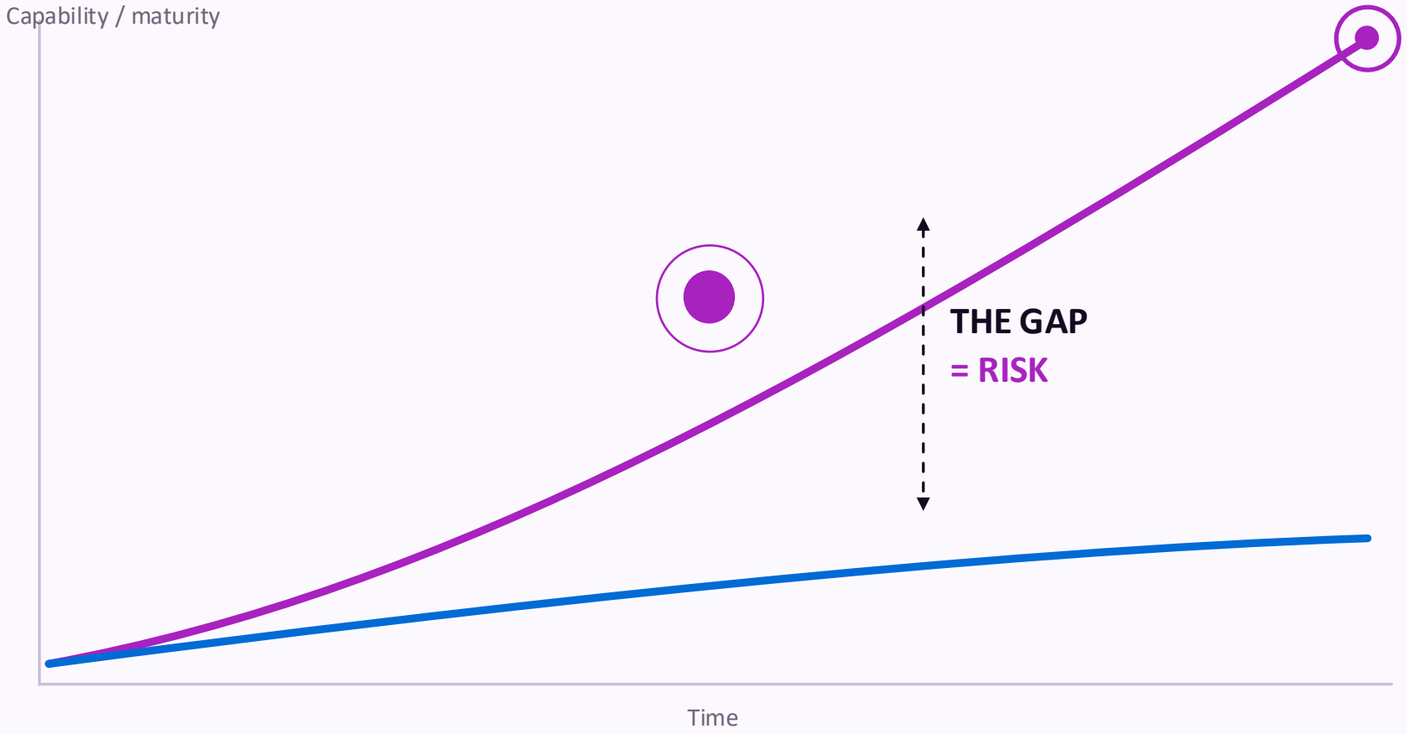

WHAT HAPPENS NEXT?

When the gap keeps widening

- Controls fail
- Assumptions don’t hold
- Exposure increases
- Impact accelerates
- Decisions can’t keep up

This isn’t a future problem. | **It’s the pressure we’re under right now.**

> **The recommendation no one seems to be focusing on: the admonishment to get the basics right.**

Kim Jones · April 2026

What Breaks First

The impact shows up here — long before anyone sees it coming.

01

Patching

→ Outages

- Emergency changes
- System instability
- Business disruption

02

Alerts

→ Noise

- Alert fatigue
- Real threats missed
- Trust erodes

03

Gaps

→ Exposure

- Blind spots
- Attack paths open
- Risk compounds

04

Identity

→ Bypass

- Privileges abused
- Lateral movement
- Defenses bypassed

These aren’t edge cases. They are where the damage begins.

Half-Prepared Isn’t Prepared

The framework is strong on prevention and governance. It’s light on response and recovery.

> “
>
> **Ready to detect and prevent at machine speed — without being ready to respond and recover at machine speed — leaves you half-prepared.**
>
> **— Summer**
>
> Left of boom is covered. The right of boom is the gap.

Four gaps to close — the right of boom

01 **IR playbooks for simultaneous, AI-speed incidents**
Three in a week, not one breach at a time.

02 **Crisis comms & regulatory notification**
Built for hour-scale timelines, not day-scale.

03 **BCPs, RTOs, and RPOs stress-tested**
Against a sub-24-hour exploitation window.

04 **Response playbooks for AI-assisted identity compromise**
Historically more damaging than zero-day exploitation.

Don’t get better at announcing bad news. Get better at recovering from it.

The Stack That Actually Matters

Security is only as strong as its foundation.

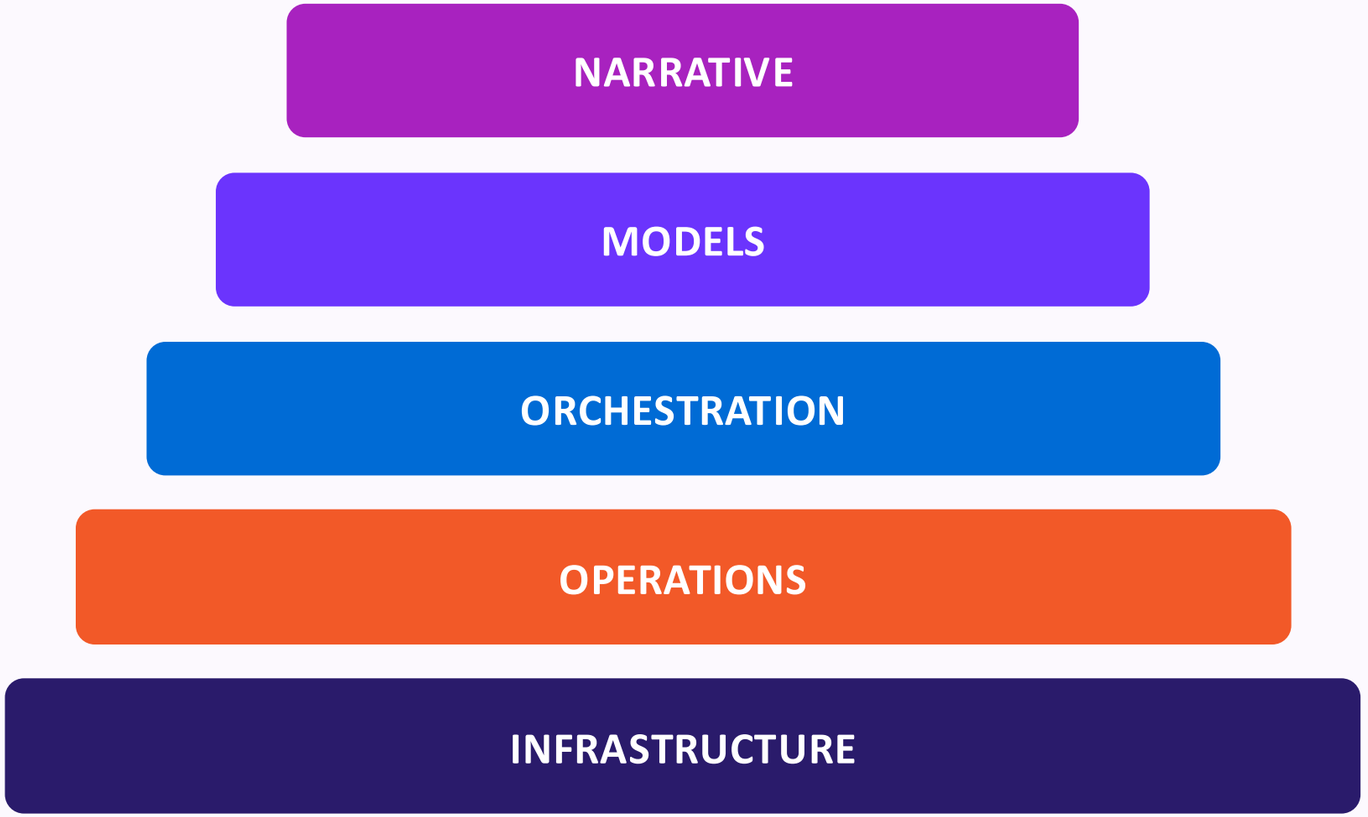

The story we tell. Shapes priorities and decisions.

The intelligence layer. AI/ML, data, and context that drive decisions.

The control plane. Workflows, integrations, and automation.

The people and processes. Execution, response, improvement.

The foundation. Compute, network, cloud, identity, and data.

You can't out-innovate a weak foundation. | **Fix the base.** | **Strengthen the stack.** | **Reduce the risk.**

The Pattern Is Now Policy

White House Executive Order · June 2, 2026 · Promoting Advanced AI Innovation and Security

What the order does

- Directs the creation of a national AI cybersecurity clearinghouse to coordinate scanning, validate vulnerabilities, and prioritize patch remediation
- Directs classified benchmarking of the cyber capability of “covered frontier models”
- Pushes AI-enabled defense and frontier-model access to critical infrastructure — rural hospitals, community banks, local utilities
- Stays voluntary — no mandatory licensing or preclearance

What it signals for your program

- The discover → validate → remediate pipeline is now a national priority
- Capability is being measured, not assumed — the curve is official
- The gap reaches the smallest operators — it’s everyone’s problem
- Innovation-first means the pace won’t slow — readiness is on you

The recognition is national. The readiness is still yours.

The future won’t wait.

Build security that can keep up.

AI is accelerating.
The gap is growing.

Risk comes from the interaction,
not the technology alone.

Security programs must evolve
from reactive to resilient.

The opportunity is here
— if we act with clarity and urgency.

This isn’t a moment.
It’s a pattern.

The question isn’t whether AI keeps accelerating.
It’s whether your systems can keep up.

Agentic AppSec Unleashed '26

by Checkmarx