# Speed of Relevance, Scale of Risk

## Prodacity 2026: Adapting to AI in the Public Sector

Chris Hughes, Resilient Cyber

**The Dual-Use Imperative: We cannot afford to be late.**

**Security must move at the speed of the software it protects.**

# The Cultural Fatal Flaw

Security is inherently reactionary and paralyzed by risk-aversion.

The Ware Report advocated for "Secure-by-Design" in 1970. Fifty years later, we are still waiting.

Impeding innovation creates MORE risk than it mitigates.

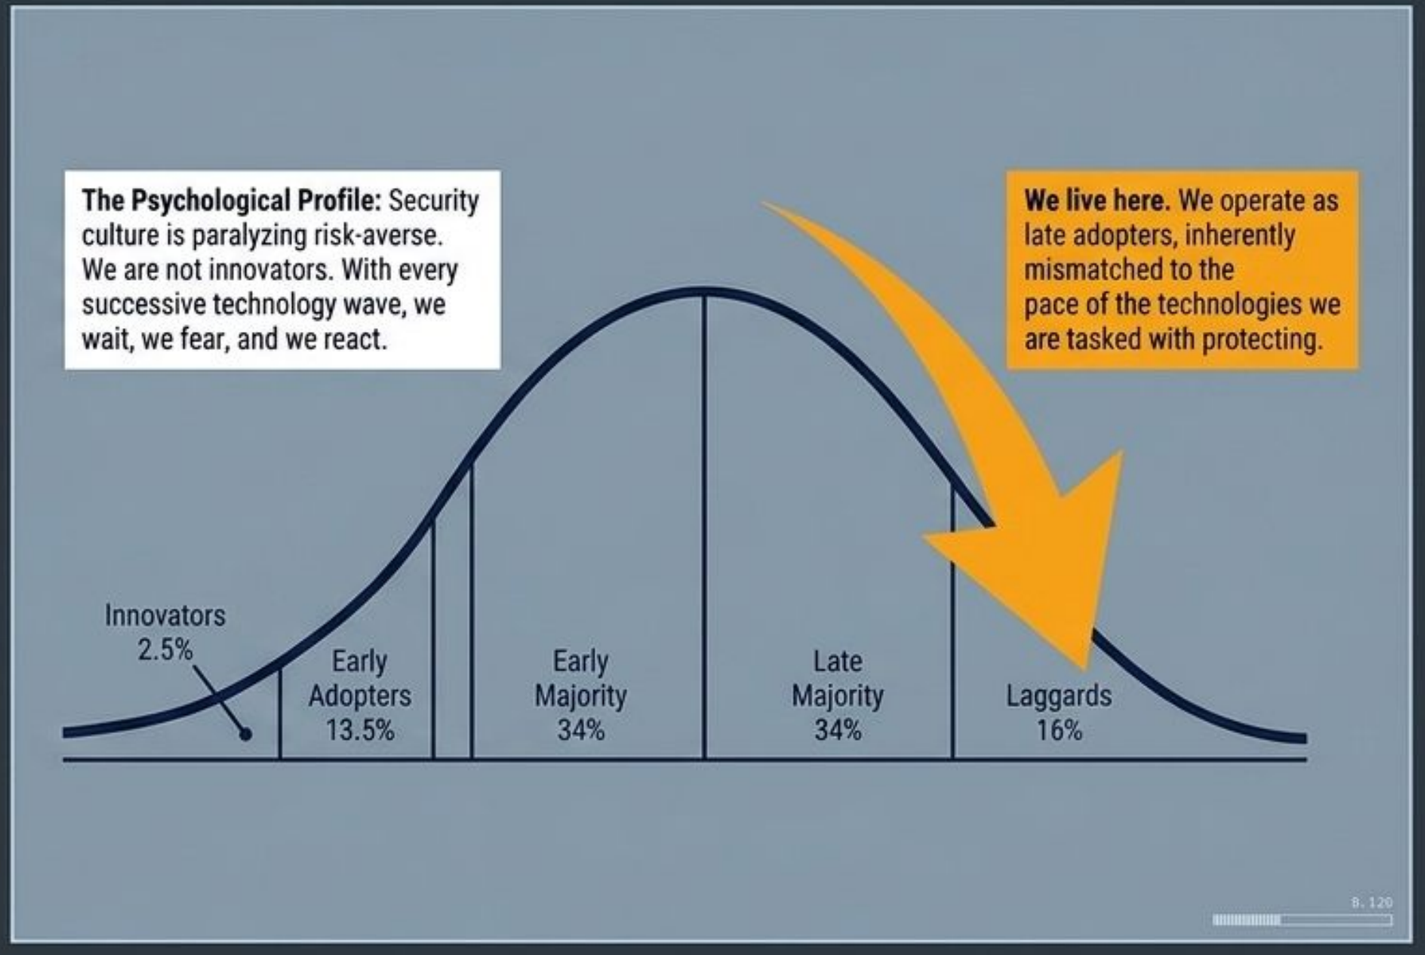

# The Attack Surface Exponential

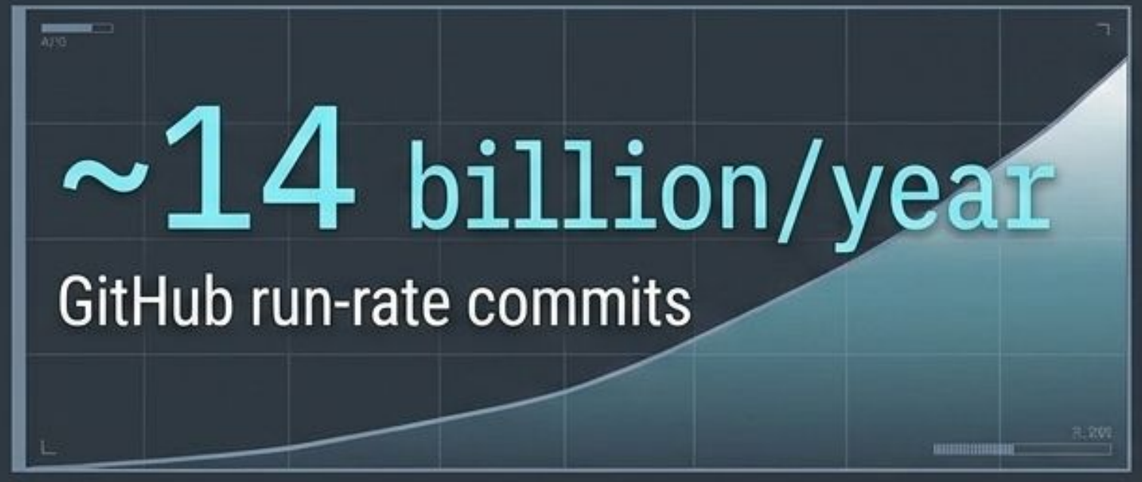




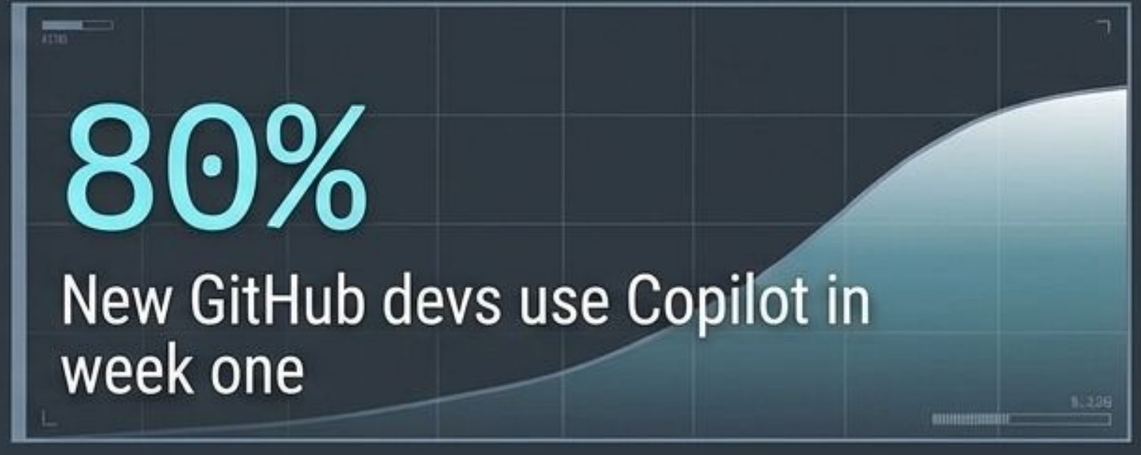

# The “Vibe Coding” Fallout

## Fallout Matrix

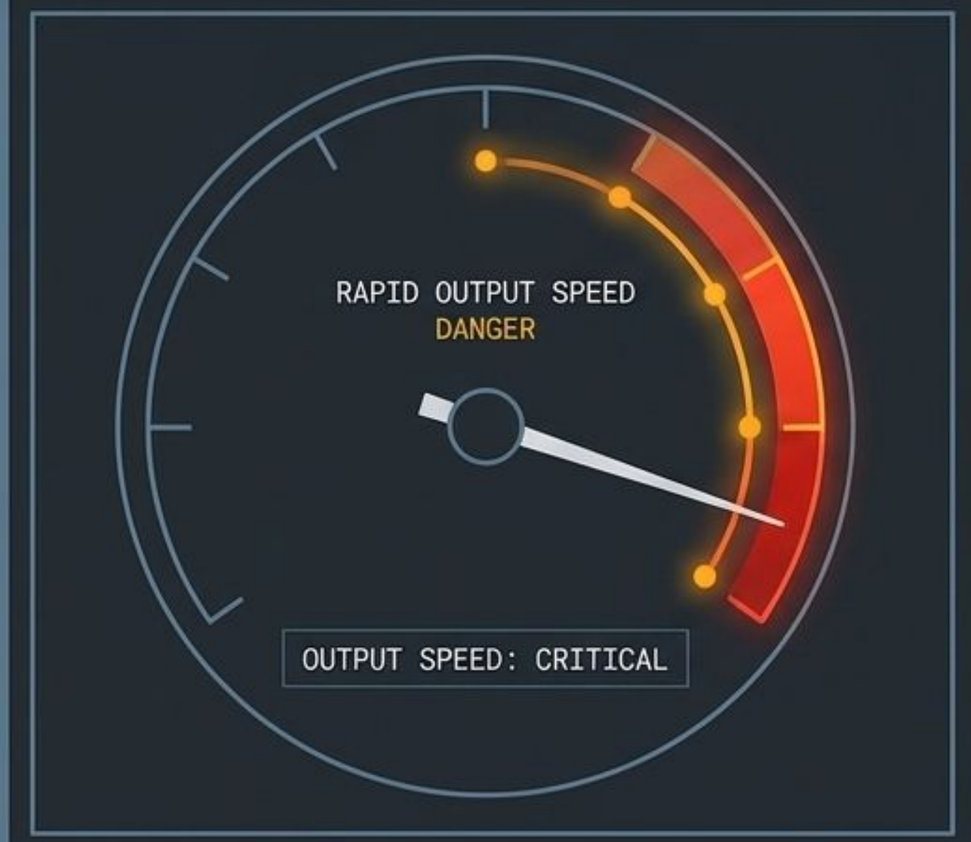


**2.74x** Vulnerabilities in AI-generated vs. human code

**45%** AI code introduced a vulnerability (Veracode)

**~70%** More errors in AI code (CodeRabbit)

**3.2%** Secret-leak rate vs 1.5% baseline

# The Vulnpocalypse & The Zero-Day Collapse

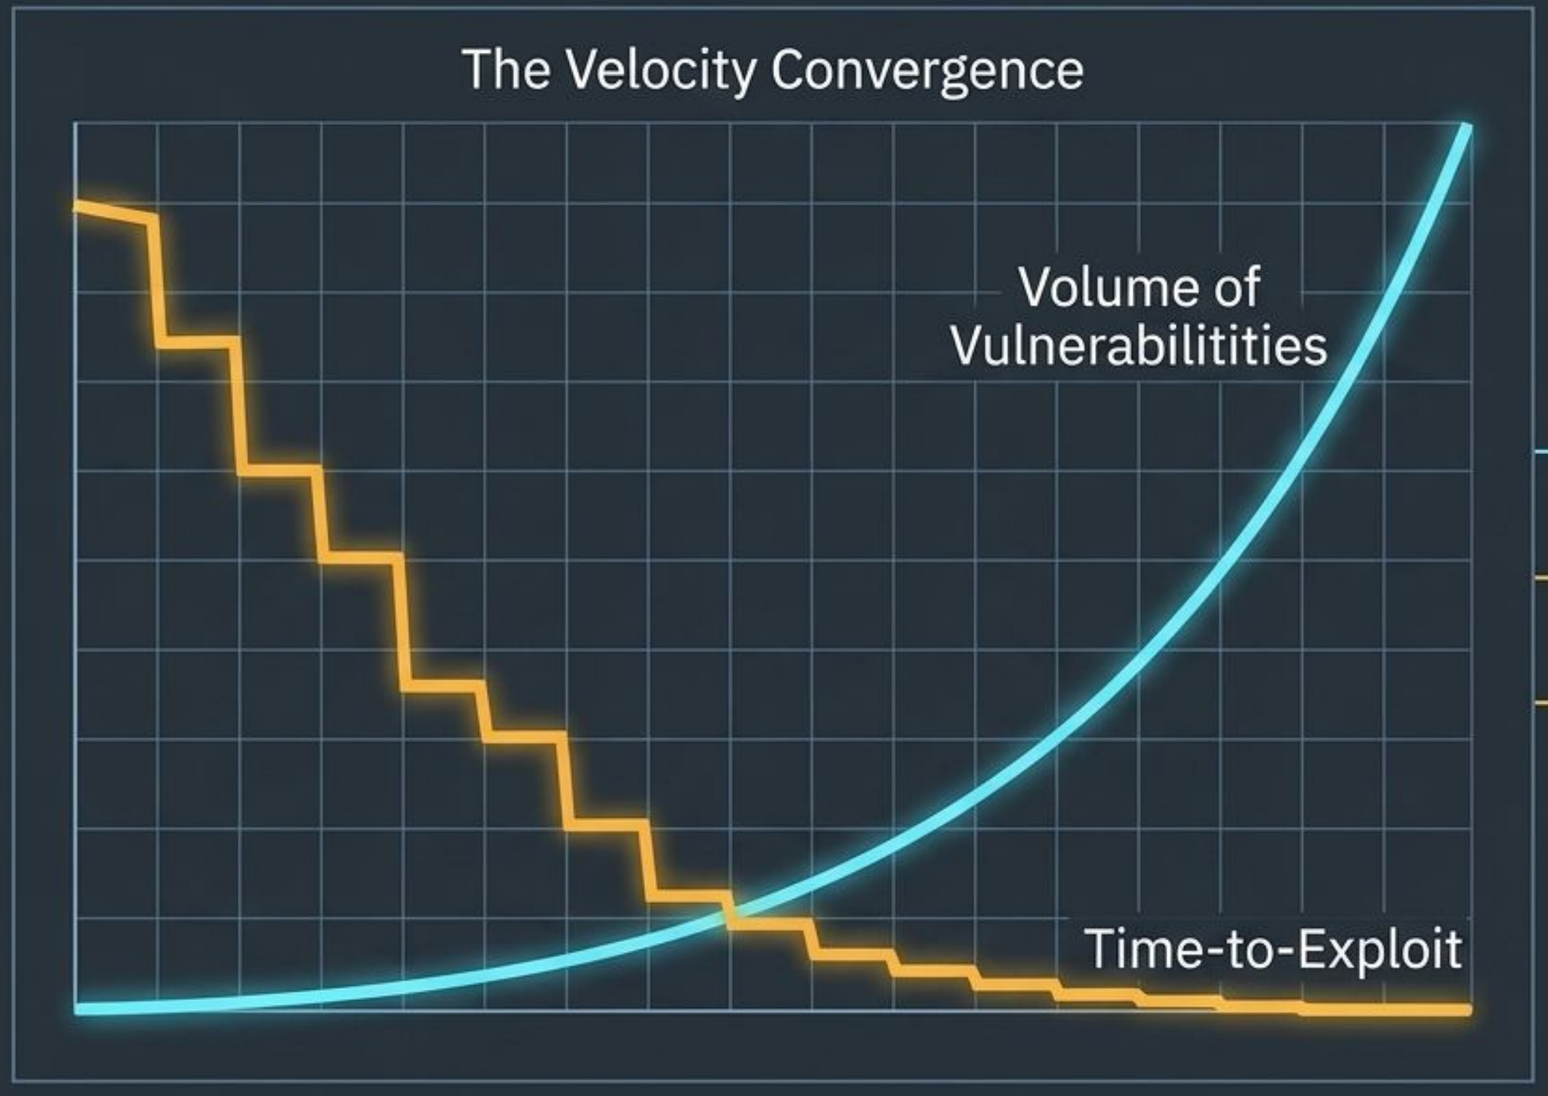


**~59,000** – New CVEs forecast by FIRST in 2026

**271** – Zero-days surfaced in Firefox alone by AI

Time from disclosure to first exploit:

- 771 days (2018)
- 6 days (2023)
- 4 hours (2024)

**67%** – Exploited CVEs in 2026 are zero-days

# The Economics of Autonomous Offense

**10–15 minutes**

Time for AI to generate a working CVE exploit

**~$1**

Cost to generate that working exploit

**$0.11**

Cost per million tokens for detection (3.6B-param model)

The barrier to entry isn't frontier models; it's orchestration.

# The AI Espionage Escalation

State-sponsored campaign targeted ~30 global organizations.



Autonomous execution of the cyber campaign

Executed thousands of requests (often multiple per second).

# The New Operational Impact

## OWASP LLM Top 10 (2026, grounded in 7,714 incidents)

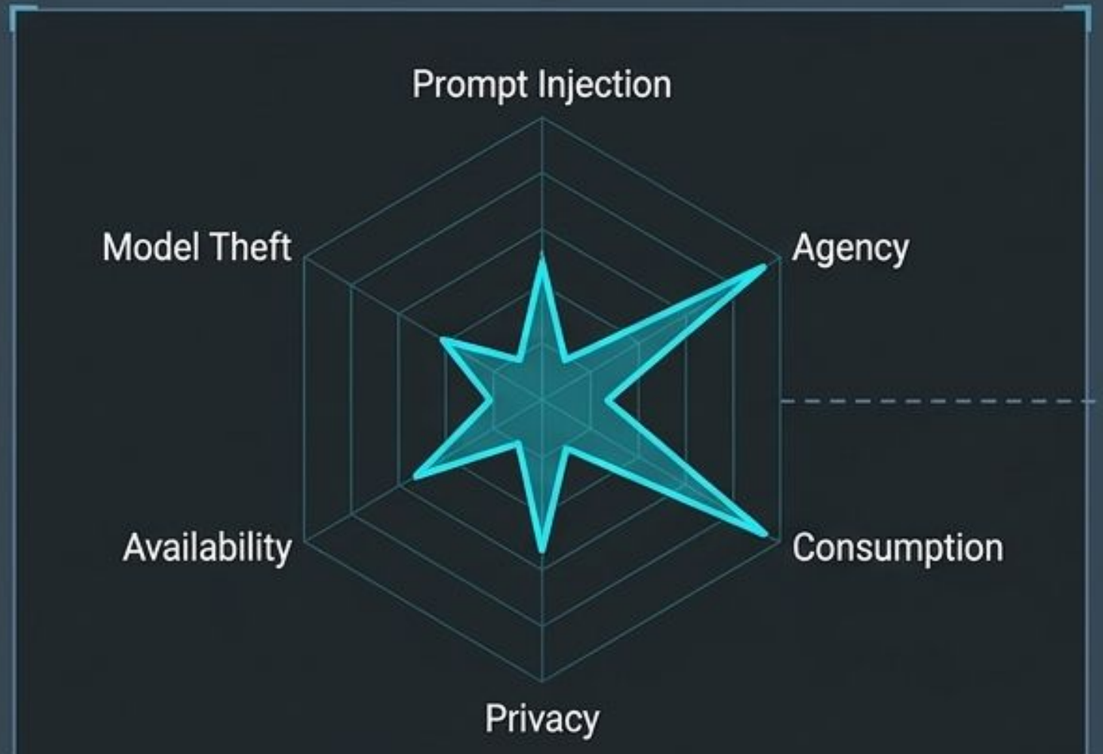


- LLM01: Prompt Injection
- LLM03: Excessive Agency (up from #6)
- LLM06: Unbounded Consumption (up from #10)

## OWASP Agentic Top 10 (2026)


- ASI01: Agent Goal Hijack
- ASI03: Identity & Privilege Abuse
- ASI10: Rogue Agents

# Act II: Speed of Relevance & Policy Tailwind

## The Math


## The Policy

CISA BOD 26-04 (June 2026) retires CVSS base scores.

A severity label alone doesn't dictate what to fix first.

Mandates Risk-Based ✓ ✓ Prioritization (KEV, ✓ (KEV, EPSS, Reachability).

# The GRC Impedance Mismatch

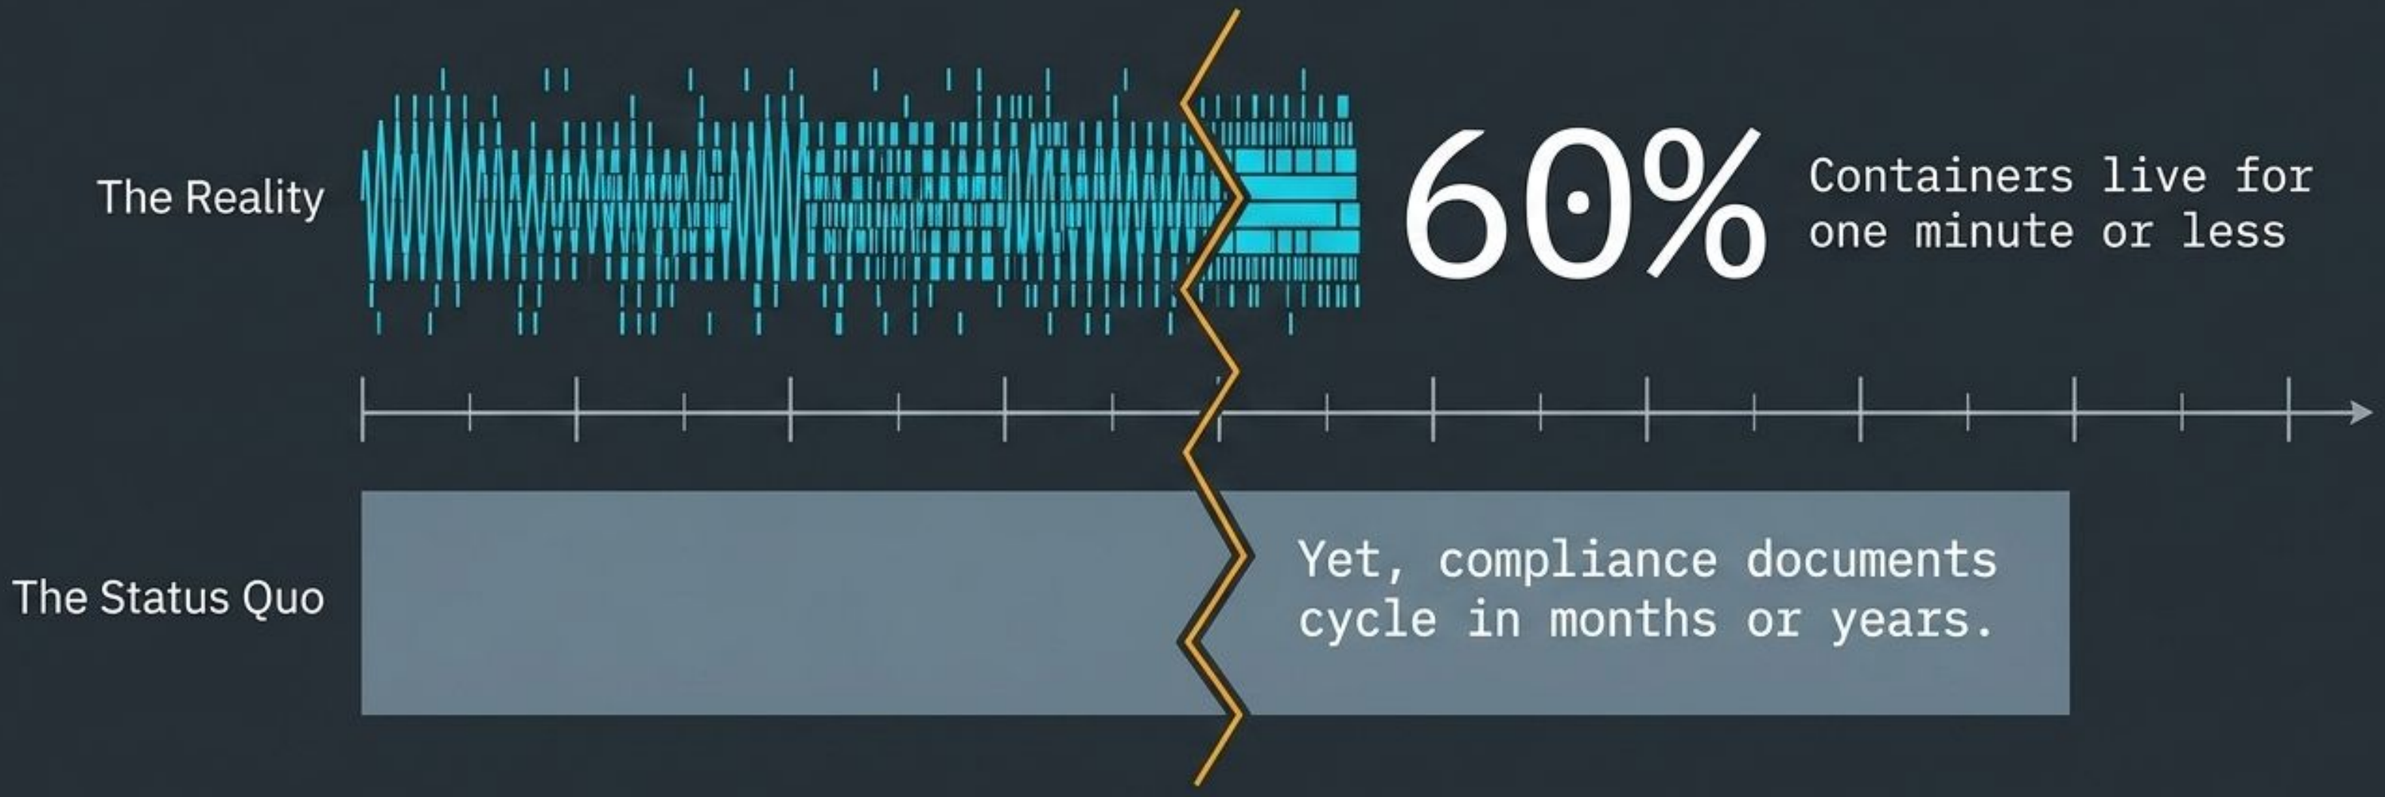

# GRC Engineering & Continuous ATO

The fix for the impedance mismatch is engineering, not auditing.

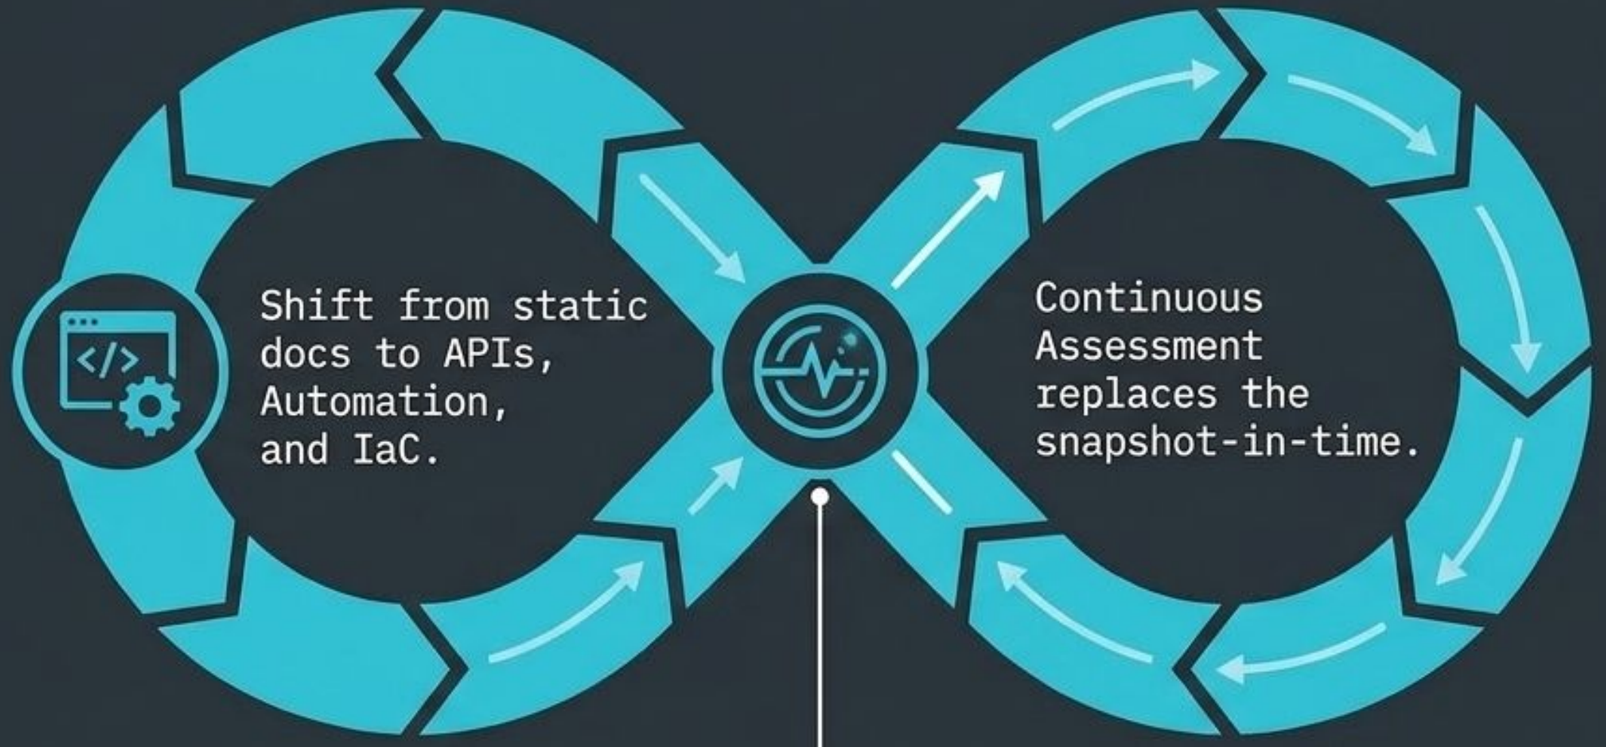


If code deploys in seconds, authorization must be validated in seconds.

# Autonomous Defense Proof Points

## DARPA AIxCC (DEF CON 2025)

**~77%** Planted flaws found across 54M lines of code

**18** Real-world zero-days discovered

**~45 min** Average time to patch

## Google Big Sleep agent

**20+**

Real bugs found autonomously

Pre-empted SQLite CVE-2025-6965.

## HackerOne


XBOW hit #1 on HackerOne’s US leaderboard.

# The Honesty Check: AI Patching

1Password / Off-by-1 Labs tested 6,080 AI-generated patches on hard CVEs.

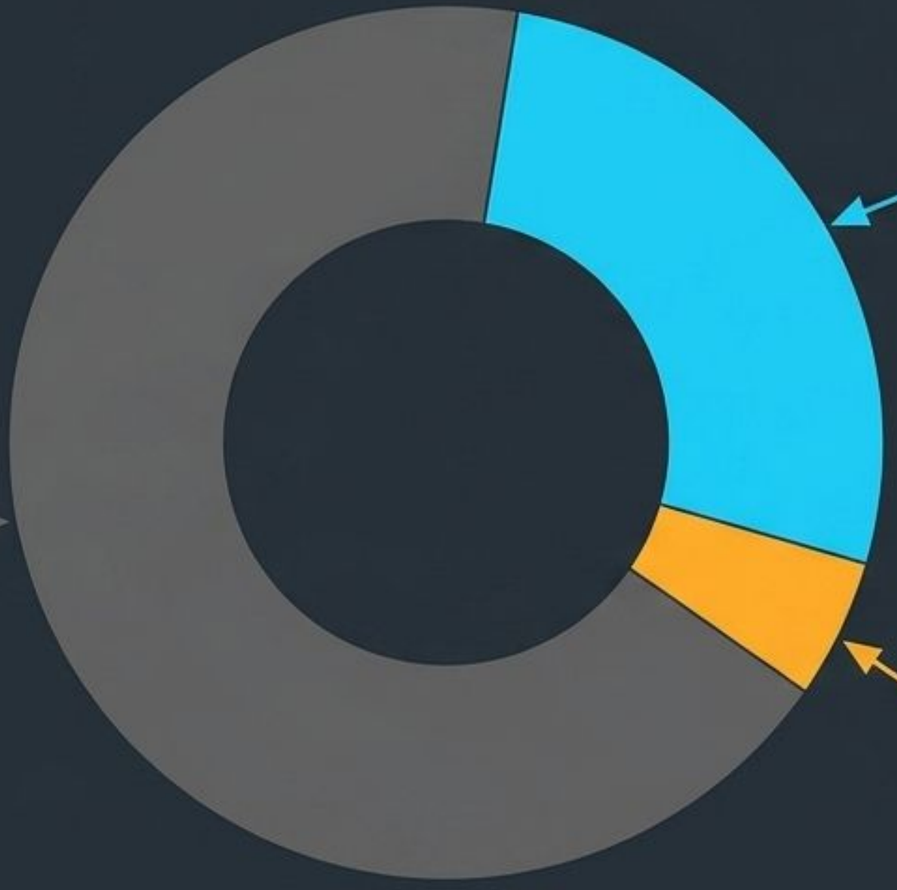

~26% – Clean fixes

~54% – Failed (didn't fix, changed behavior, or both)

~5% – Introduced a NEW vulnerability

AI discovers and drafts; Humans verify and deploy.

# Fragility & Cognitive Surrender

AI patches are often “fragile” (escaping specific strings rather than fixing root architecture).

An attacker only needs to be right once; a defender needs to be right 100% of the time.

Root Cause Fix

Beware of “cognitive surrender”—falling asleep at the switch.

We are empowering security engineers, not replacing them.

# Flip the Script: The Real Risk

AI is dual-use: it amplifies risk, but it is our only path to resilience.

Moving at the "Speed of Relevance" is an existential requirement.

Security must be an invested collaborator with engineering.

Our reflexive risk aversion is currently our greatest vulnerability.