AI & The Vulnerability Lifecycle

Has AI **actually** changed the zero-day curve?

SOURCES GTIG · Mandiant M-Trends 2026 · Rapid7 · AISI · IBM · Veracode · Edgescan · Vulncheck · As of May 2026

01 FRONTIER MODEL CAPABILITY

~100%

on CyberBench 4.6
Claude Opus 4.6 reached the ceiling of Anthropic's cybersecurity benchmark, signaling a need for more advanced testing to assess future capability gains.

Anthropic System Card, Feb 2026

90.5%

GPT-5.5 on narrow cyber tasks (pass@5) —AISI

UK AI Safety Institute

$2.77

Average cost per CVE reproduced by AI (CVEGenie, 51% success rate)

arXiv, Sept 2025

6/10

Claude Mythos completes a full 32-step enterprise attack chain—recon through domain takeover—6 out of 10 attempts. GPT-5.5 does 3/10.

AISI Cyber Range Eval

! Anthropic publicly admitted its benchmark stack **can no longer ceiling-check offensive cyber capability**—the first vendor acknowledgment of this kind.

02 ZERO-DAY LANDSCAPE—SHAPE CHANGED, VOLUME DIDN'T

IN-THE-WILD ZERO DAYS EXPLOITED

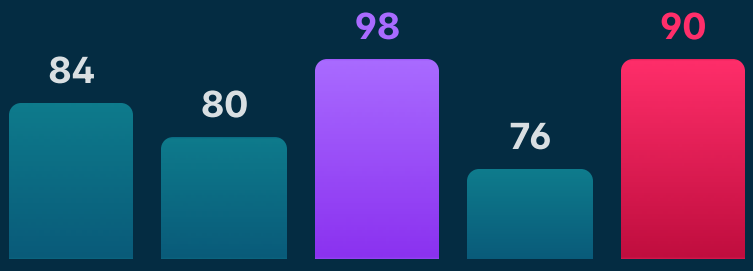

▲ Rebound, not surge

48% of 2025 zero-days targeted **enterprise products**—an all-time high (VPNs, SAP, SharePoint, Ivanti)

18 vs 15 Commercial **spyware vendors** out-attributed nation-states for the first time (CVs vs. state actors)

<0.5% of ~46,000 2025 CVEs publicly credit **AI-assisted discovery** (Barracuda Mythos Hype Index)

03 EXPLOIT TIMELINE—EXPLOITATION NOW PRECEDES PATCHING

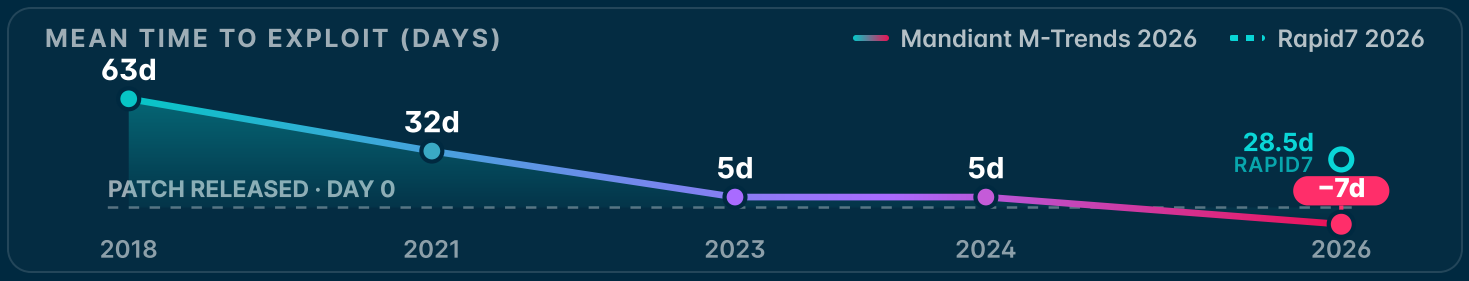

−7d exposure window. On average, the worst CVEs are exploited seven days **before** a patch is available or publicly disclosed — making the exploitation window effectively "negative."

29% of 2025 KEVs exploited **on or before** CVE publication (VulnCheck State of Exploitation 2026)

8× growth in **edge-device** exploitation as a share of all exploit actions YoY (Verizon DBIR 2025)

22 sec median **access-broker handoff** to ransomware affiliate (2025, Mandiant)

04 THE DEFENDER GAP—REMEDIATION ISN'T KEEPING PACE

MEAN TIME TO REMEDIATE CRITICAL VULNS (DAYS) · EDGESCAN 2024

Application / API layer **74.3d**

Device / network layer **54.8d**

Average days to remediate critical & high-severity vulnerabilities, by stack layer.

EXPOSURE & PATCH BACKLOG

Orgs carrying security debt **82%**
Veracode 2026

Critical KEV still open at day 7 **63%**
Qualys TRU

Need 1 week+ to deploy patches—against a 48-minute attacker breakout window **77%**
ABDITIVA survey

$1.9M saved per breach w/ **extensive AI & automation** (IBM 2025)

241 days mean time to identify + contain breaches—**9-yr low** (IBM Cost of a Data Breach 2025)

63% of breached orgs have **zero AI governance** policy (IBM Cost of a Data Breach 2025)

+47% fix-time growth since 2020 **(Veracode —'Drowning in Security Debt')**

05 SAFEBREACH 2026 STATE OF THE BREACH—1.8M SIMULATIONS

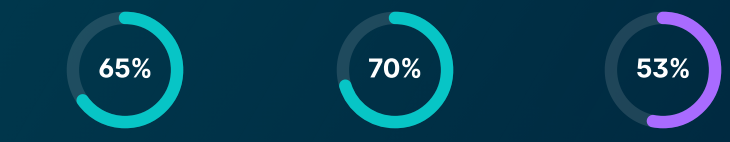

SafeBreach 2026 State of the Breach: 1.8 million attack simulations across customer enterprises reveal actual control effectiveness—not hypothetical. Validate continuously.