cymulate

# The AI Exposure Gap

How agentic CTEM closes the distance between attacker speed and defender response

## INTRODUCTION

**AI has changed the math of exposure management.**
Attackers now find, weaponize and exploit weaknesses at machine speed. Defenders have to respond the same way.



Frontier AI models can audit code, reverse engineer binaries and assist exploit development at a scale no human team can match. At the same time, every AI assistant, agent and LLM-powered application an organization deploys adds new identities, integrations and data paths for attackers to target.

Continuous threat exposure management (CTEM) was designed to help security teams see their environment from an attacker's perspective and act before exposures are exploited. Yet many CTEM programs still run on manual research, theoretical prioritization and patch queues measured in weeks.

This e-book explores how AI is reshaping threat exposure, why first-generation CTEM programs stall and how security teams can use AI, automation and continuous validation to **prove** what's exploitable, **prioritize** what matters and **adapt** their defenses before attackers strike.

WHAT'S INSIDE

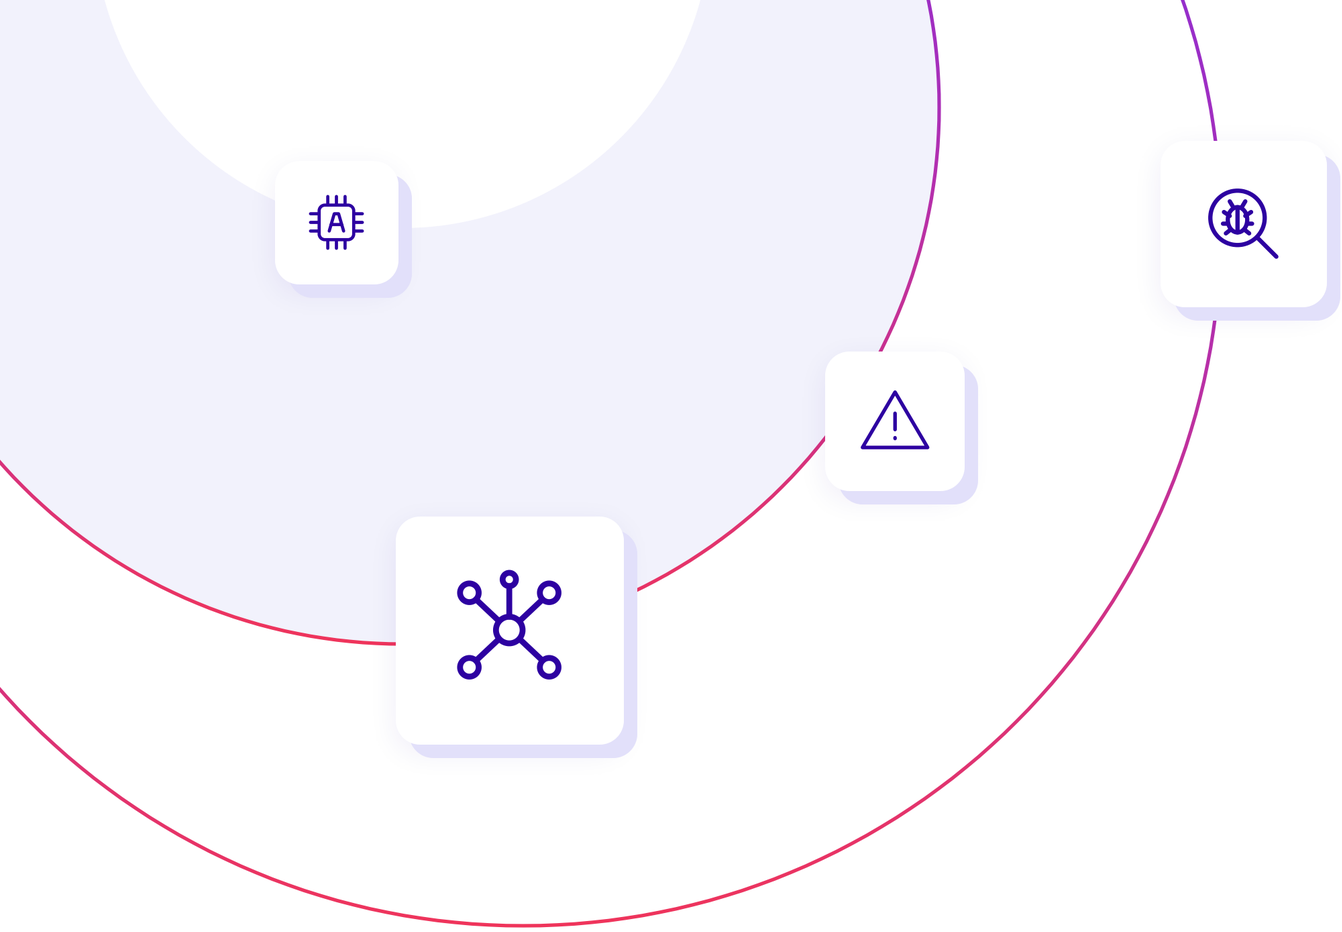
# From AI-Driven Exposure to Exposure-Informed Defense


01 **AI Is Rewriting the Exposure Equation**
Why AI is a new source of threat exposure

02 **CTEM Starts With an Outcome**
Where AI accelerates each phase of the framework

03 **Four CTEM Fails in the AI Era**
And how to fix them

04 **Agentic CTEM: Adapt at Machine Speed**
Autonomous routines and exposure-informed defenses



01

# AI Is Rewriting the Exposure Equation

**~50% more exposures**
**expected to be discovered in 2026 than in 2025**

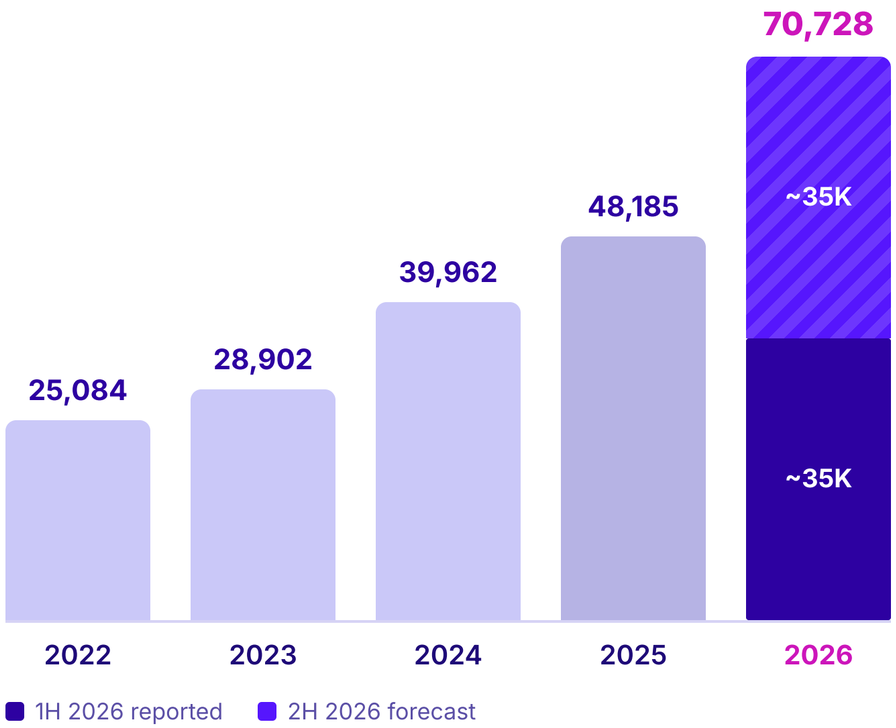
Long before frontier AI models made headlines, vulnerability management teams were already drowning in more vulnerabilities than any team could manage with speed and scale. AI is turning that flood into a surge.

AI models can continuously audit codebases, reverse engineer binaries and accelerate exploit development. The result: more exposures discovered, faster weaponization and a window between disclosure and attack that is shrinking from weeks to hours – and sometimes minutes.

The danger isn't a single AI-generated exploit. It's the growing gap between attacker speed and the slow, fragmented processes most organizations still use to reduce exposure.

## Published CVEs per year

Disclosed vulnerabilities have nearly tripled since 2022, with AI-assisted discovery driving the 2026 surge.




*Source: Cymulate, CTEM at Machine Speed webinar, 2026. 2026 is a full-year estimate.*

**~25%**

of known exploited vulnerabilities (KEVs) are exploited within 24 hours of disclosure

**32%**

of SecOps teams say they have too many exposures to prioritize effectively

*Source: Cymulate Threat Exposure Validation Impact Report 2025*

## Once inside, AI-assisted adversaries move in minutes – or seconds

**27 seconds**

fastest observed adversary breakout time

**29 minutes**

average adversary breakout time

*Source: CrowdStrike*

A remediation process measured in days or weeks can’t keep pace with an attack that unfolds in minutes.

AI CUTS BOTH WAYS

# AI Is a New Source of Threat Exposure

AI creates risk in two directions at once. It gives attackers speed and scale, and it gives every organization adopting it a new class of assets to scope, discover and validate.

## AI accelerates the attacker

- Automated vulnerability discovery across source code and binaries
- Faster exploit development and weaponization
- Attack path analysis that chains minor exposures into major breaches
- Targeted campaigns launched at scale in hours, not weeks

## AI expands the attack surface

- AI assistants, copilots and desktop AI with broad access to data
- AI agents with identities, API keys and privileges to govern
- LLM-powered apps exposed to prompt injection and data leakage
- New integrations and data pipelines outside traditional scanning coverage

**Both sides raise the stakes for CTEM:**

1. **More exposures to manage**
2. **Less time to act and new asset classes that must be in scope**

*Frontier AI models are increasing both the volume of exposure insights available to defenders and the speed at which threat actors identify, research and potentially exploit weaknesses. As a result, organizations must improve how they prioritize, coordinate and mobilize exposure reduction activities.*

Gartner, Shape Your CTEM Operating Model for Frontier AI Models and the Agentic Era, Mitchell Schneider, Jonathan Nunez, November 6, 2026.

02

# CTEM Starts With an Outcome

Gartner® created the CTEM framework to give security teams a continuous way to see their environment through an attacker's eyes and reduce risk before exposures are exploited.

Never intended to be a simple evolution of traditional vulnerability management, CTEM was designed to enable the whole security organization with clarity, alignment and action.

### Clarity
A clear view of real security gaps

### Alignment
Shared priorities and an actionable plan across teams

### Action
Hardened security before exposures are exploited

A NEW PERSPECTIVE

**CTEM reorients priorities around three questions**

**01**
**How do we look from an attacker's point of view?**

**02**
**Which configurations leave us exposed?**

**03**
**How would our controls and response cope?**

# The CTEM Framework, Accelerated by AI

Each of the five CTEM phases still matters. What changes in the AI era is how fast each one must run, and how much of the repeatable work AI and automation can take on.



## Scope

Focus on the assets and business services that matter most – including AI apps and agents

**WHERE AI HELPS**

Environment-aware AI agents baseline posture and business context



## Discover

Every source of exposure, not just CVEs: misconfigurations, identities, cloud paths and control gaps

**WHERE AI HELPS**

Correlates exposures to threats across integrated discovery tools



## Prioritize

Stack-rank exposures by threat context, asset importance and business impact

**WHERE AI HELPS**

Ranks by proven exploitability instead of modeled theory



## Validate

Prove what an attacker can exploit and whether controls prevent or detect the attack

**WHERE AI HELPS**

Agentic AI tailors and runs safe attack simulations against the latest threats



## Mobilize

Coordinate teams to reduce exposure through the fastest effective action

**WHERE AI HELPS**

Automates control updates and recommends the next best action

03

# Four CTEM Fails in the AI Era

Many first-generation CTEM programs have the right framework but the wrong operating model. Old vulnerability management habits are hard to break, and AI-powered discovery, weaponization and targeted campaigns punish every delay.

Four failures show up again and again.

**FAIL #1**

**Validation without evidence of control effectiveness**

**FAIL #2**

**Patching treated as the only response**

**FAIL #3**

**Vulnerability management with a new name**

**FAIL #4**

**Collaboration that creates slow, manual handoffs**

*Your objective is to shift exposure management from counting vulnerabilities to measuring how quickly attacker-relevant exposures are identified, validated and mitigated within defined response objectives.*

Gartner, Modernize Exposure Management to Counter AI-Augmented Vulnerability Discovery, Mitchell Schneider, Jonathan Nunez, September 24, 2026.

CTEM FAIL #1

# Validation Without Evidence

Many CTEM programs treat passive analysis as validation. Attack surface management maps how an attacker could theoretically reach an exposure. Security control configuration audits model which threats a control should stop. Both are useful. Neither proves what happens when a real attack runs in your environment.

A theoretical attack path doesn't show whether an endpoint policy will stop execution, a network control will block lateral movement or a SIEM rule will fire an alert. A path that looks urgent may already be blocked. A moderate-rated exposure may lead straight to a critical system because a detection rule is missing.

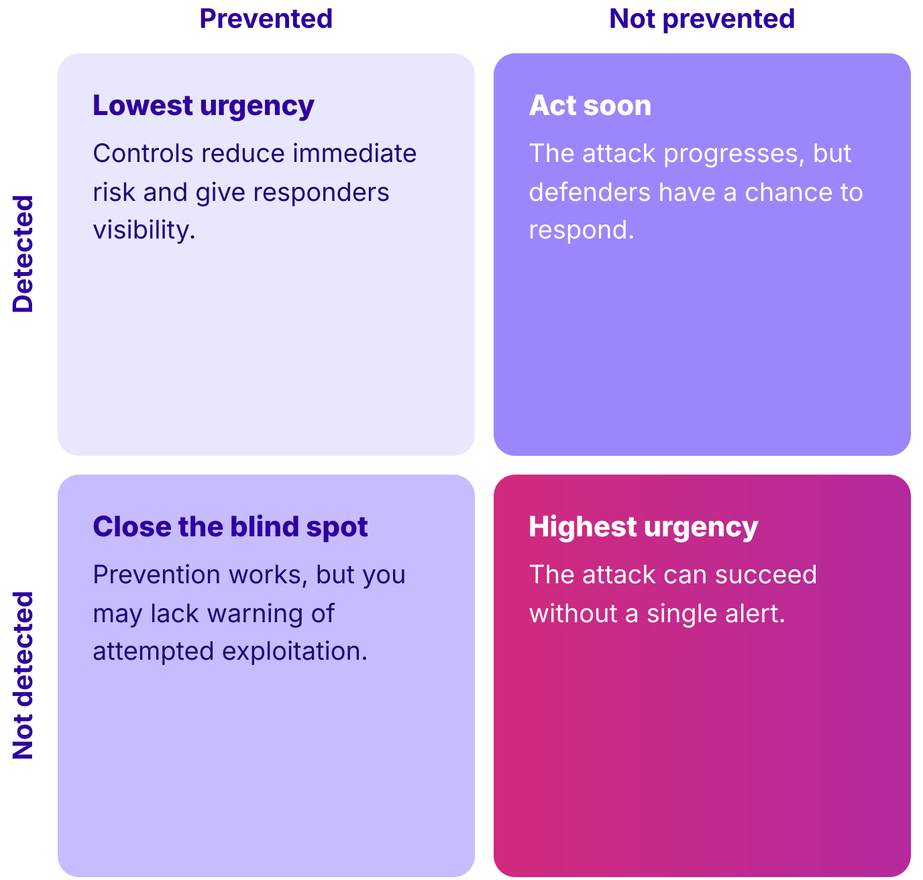
THE FIX

### Make validation the central filter

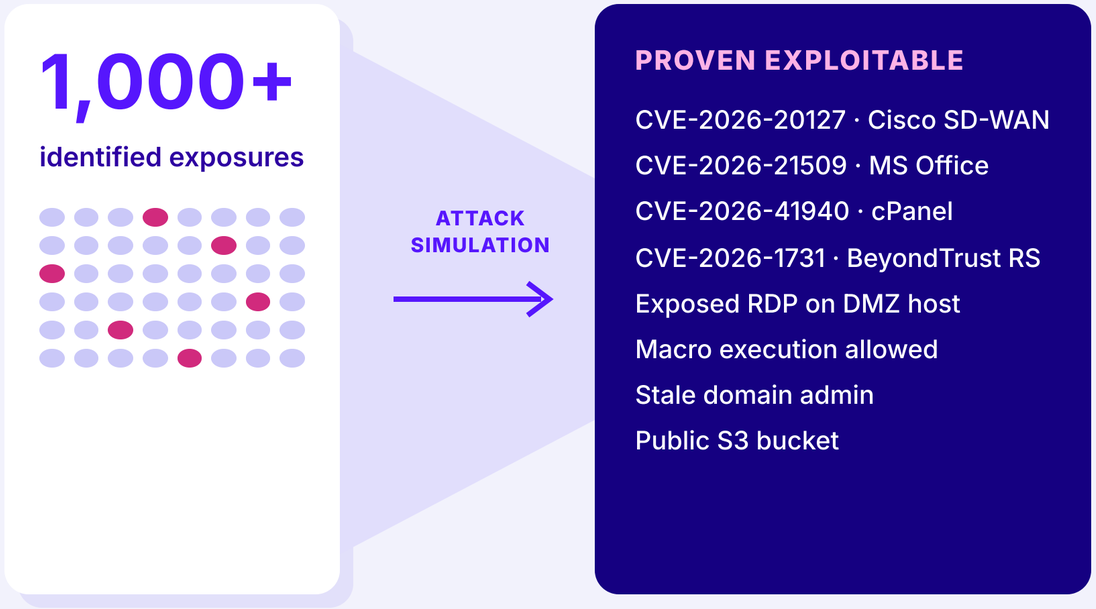
Use automated attack simulation to safely reproduce the attacker behavior tied to each exposure. Measure whether controls prevent it, detect it and trigger the expected response.

## Four validation outcomes

Every team sees the same evidence – and the same urgency.



# Validation Turns Thousands of Exposures Into a Focused List




*Illustrative example of validated exposures.*

*"Cymulate helps us evaluate whether a vulnerability is exploitable in our environment and enables us to prioritize remediation efforts."*

**– Adam Champion, Head of Information Security, Saffron Building Society**

## Validation has to run at machine speed

When AI compresses attacker timelines, validation can't wait for a quarterly pen test or a manual research cycle. AI-driven validation maps a new threat or CVE to your environment, builds a tailored attack scenario and runs it safely – often before an analyst opens the ticket.

The evidence gives every team the same facts: which attacks reached their target, which controls failed, which behaviors went undetected and which changes reduced risk on retest.

**NEW KPI**

**Time to validate exploit (TTVE):** how long it takes to confirm whether a new exposure is exploitable in your environment.

CTEM FAIL #2

# Patching Treated as the Only Response

Patching remains essential, but it rarely moves at the speed of exploitation. A patch may need testing, a maintenance window, owner approval and a rollback plan. Some exposures have no vendor patch at all. When AI shrinks the window between disclosure and exploitation, a patch-only program inherits every one of those speed limits.

THE FIX

### Mobilize the fastest effective mitigation first

Ask one question: which available action reduces risk fastest? Block indicators of compromise, tune a detection, apply a WAF rule or virtual patch, or restrict privileges. Then retest against the same attack and track **time to protection** alongside time to patch.

### Time to build, deploy and test in production

Automation compresses every mitigation option.

| | *Manual* | *Automated* |
|---|---|---|
| **IoC block**<br>Across all controls | **1–2 hrs.** | **< 2 min.** |
| **New behavioral rule** | **2–10 days** | **30 min.** |
| **Control configuration** | **1–5 days** | **1–8 hrs.** |
| **Patching** | **2–60 days** | **1–3 days** |

*Source: Cymulate, CTEM at Machine Speed webinar, 2026.*

*By year-end 2027, cybersecurity solutions will also use AI to improve the ability to deploy and test compensating controls and produce detailed attack path modeling, reducing the risk of a breach by 40%.*

Gartner, Prioritize and Compensate: Exposure Management for the AI Era, Craig Lawson, Mitchell Schneider, Chris Saunderson, October 14, 2026.

CTEM FAIL #3

# Vulnerability Management With a New Name

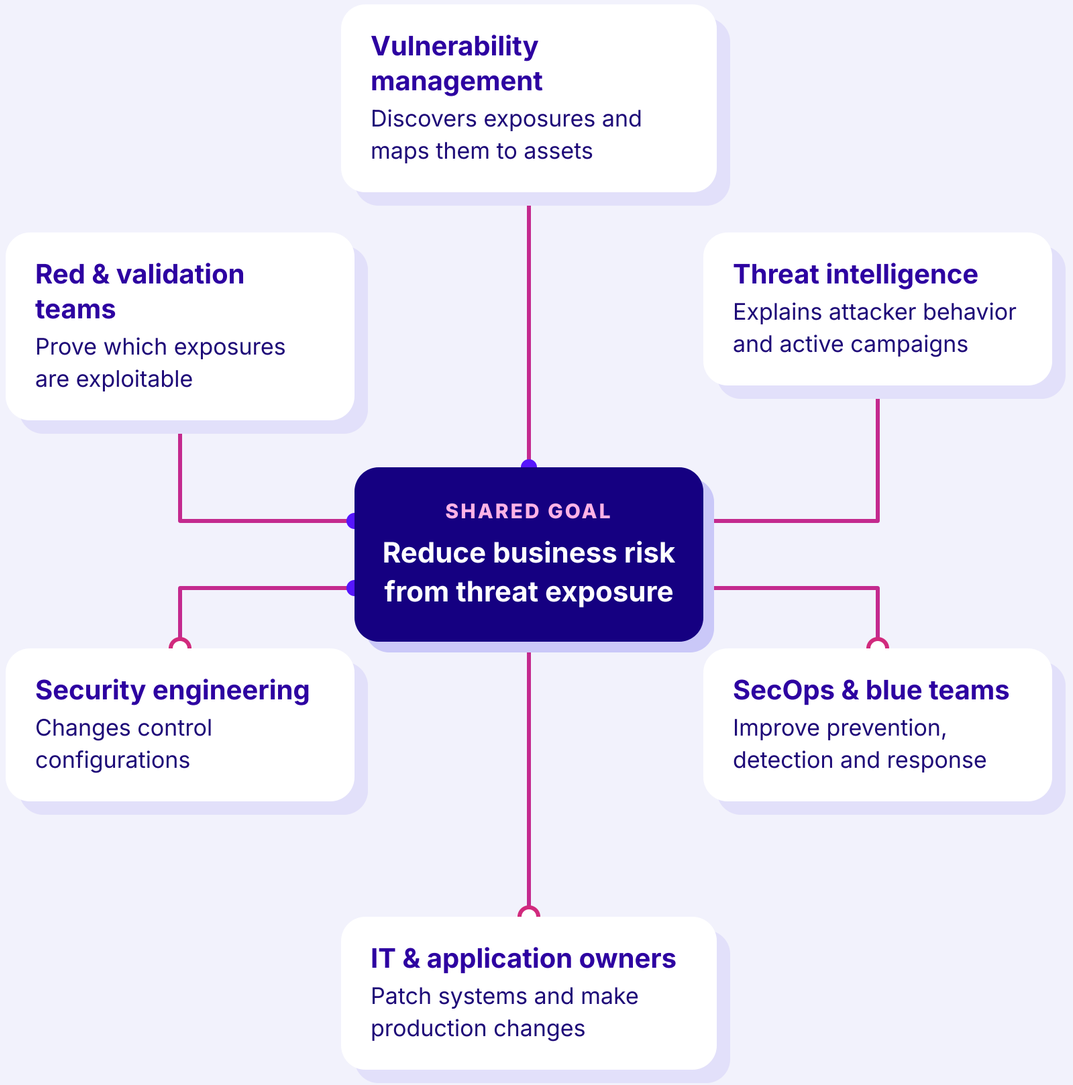
For many organizations, CTEM is a slow evolution of vulnerability management: the same process with added asset context, threat intelligence and attack-path data. That's progress, but it isn't CTEM if the work stays inside one team.

A CVE-centered view also misses misconfigurations, excessive privileges, missing detections – and the new identities and integrations that AI agents bring with them.

THE FIX

**Run CTEM as a program-wide operating model**

Define roles, ownership and escalation across teams, then share one set of metrics: validated gaps, control failures, mitigation status, time to protection and proof that each fix worked.



CTEM FAIL #4

# Collaboration That Creates Slow, Manual Handoffs

Moving CTEM beyond vulnerability management brings more teams in – each with its own systems and priorities. A single exposure may pass from a scanner to vulnerability management, threat intelligence, the red team, SecOps and IT. Every handoff means interpreting someone else's work, opening a ticket and waiting. A process measured in days can't answer an AI-assisted attack that unfolds in hours.

> *Closing this gap requires a shift from technology-centric, severity-based approaches to time-based, outcome-driven exposure management, supported by organizational change and end-to-end automation rather than point solutions.*
>
> Gartner, Modernize Exposure Management to Counter AI-Augmented Vulnerability Discovery, Mitchell Schneider, Jonathan Nunez, September 24, 2026.

## THE FIX Automate the workflow from discovery through retesting




**Humans stay in control.** AI agents handle the research, coordination and repeatable execution. People review the evidence and approve sensitive changes.

04

# Agentic CTEM: Adapt at Machine Speed

AI created the speed problem. AI and automation are also the only practical way to match it.

An agentic CTEM system moves continuously through six connected steps. AI agents read threat intelligence, map a new threat to the environment, build a test, run it and recommend the next action. Automation deploys the approved change and runs the test again.

**By the time an analyst reviews the exposure, the team may already know whether the threat is relevant, which controls failed and which mitigation worked.**



# From Days to Hours: Agentic Routines in Action

Cymulate Exposure Validation provides the automation that turns days of work into hours. Powered by Vero AI, Cymulate Cowork assemble squads of agents and skills to execute autonomous routines that operate CTEM at machine speed.

| | BY HAND | CYMULATE AUTOMATION | CYMULATE COWORK |
|---|---|---|---|
| **A new threat in the news** | read · map · script · run<br>**Days** | build & launch<br>**~1 hour** | agentic routine<br>**VALIDATED** |
| **A critical CVE drops** | reproduce · patch blind<br>**Days** | find scenario · launch<br>**~1 hour** | agentic routine<br>**MITIGATED** |
| **A detection gap** | hand-write a rule · hope<br>**Days** | replay · tune<br>**1–2 hours** | agentic routine<br>**DETECTION DEPLOYED** |
| **The board asks** | chase · reconcile · build<br>**Weeks** | export · assemble<br>**1–2 hours** | agentic routine<br>**REPORT CREATED** |

**Cymulate Exposure Validation automation turns days into hours. Cymulate Cowork does the work for you.**

# How Cymulate Builds Exposure-Informed Defenses

Cymulate, powered with Vero AI, connects exposure and threat context to continuous attack validation and control optimization – then pushes the results back into your security controls.



## THE BOTTOM LINE:

AI is increasing the speed and volume of threat exposure, and it's adding new AI assets that attackers will target. CTEM programs built on manual research, theoretical prioritization and patching alone can't keep pace.

The answer is to give defenders the same advantage. Use AI and automation to validate what's exploitable in your environment, prioritize with evidence and mobilize the fastest effective mitigation – then retest to prove the risk is gone.

### Key takeaways

- Run CTEM as a program-wide initiative, with AI apps and agents in scope
- Validate to focus immediate action on what's exploitable for you
- Invest in AI and automation that remove handoffs across teams
- Mobilize at machine speed to build exposure-informed defenses

**MEASURE WHAT MATTERS**

| | | | | |
|---|---|---|---|---|
| **Time to validate exploit (TTVE)** | **Time to protection** | **Validated gaps closed** | **MTTD and MTTR** | **MITRE ATT&CK coverage** |

cymulate

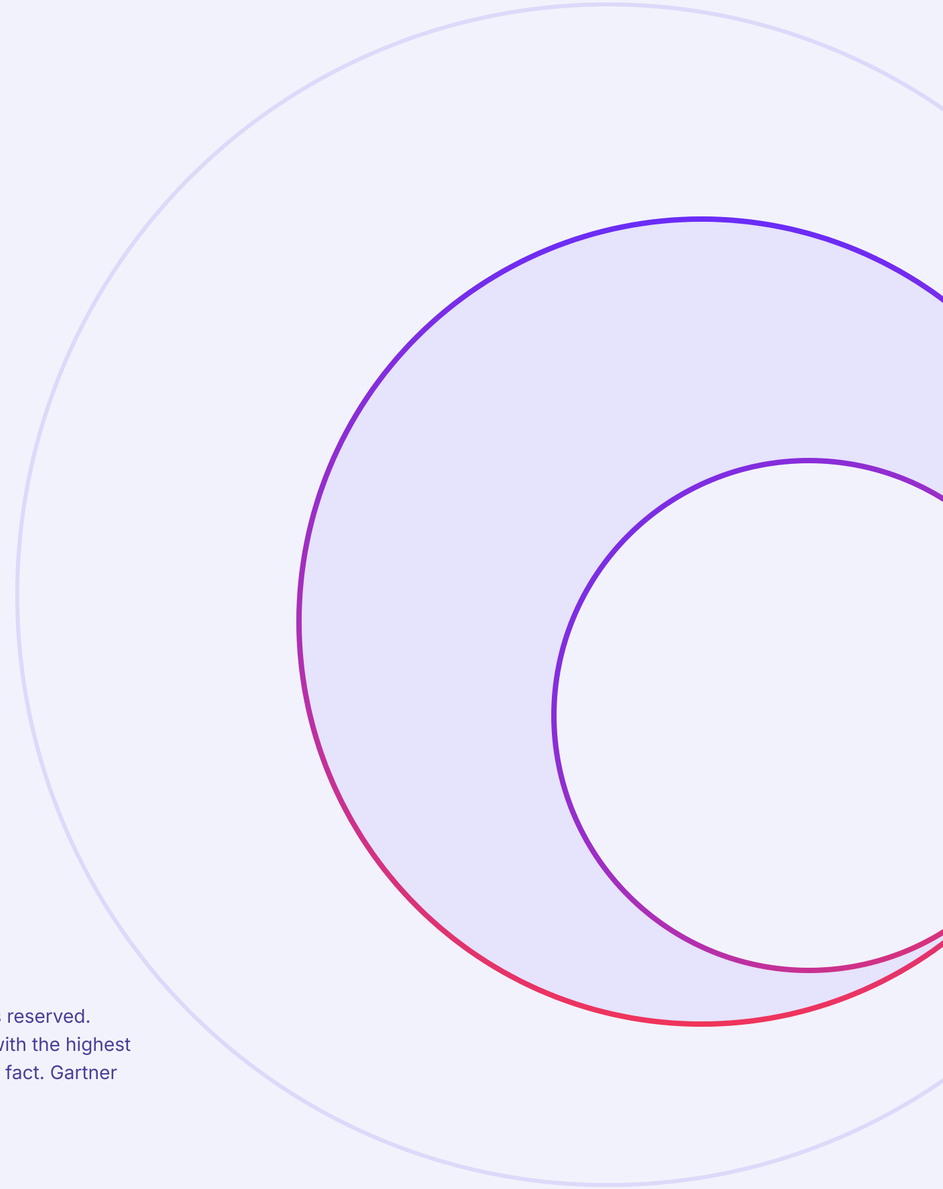
## See CTEM at machine speed

Get a personalized demo of how Cymulate applies AI and continuous validation to build exposure-informed defenses.

**Book a Demo**

## About Cymulate

Cymulate is the leader in proactive, AI-powered security that continuously proves, prioritizes and adapts against real attacker behavior – before incidents occur. More than 1,000 enterprise security teams rely on Cymulate for autonomous threat validation and cyber defense engineering. Founded and led by experienced red teamers who know that testing alone does not deliver better security, Cymulate goes beyond threat validation to build threat- and exposure-informed cyber defenses. For more information, visit www.cymulate.com.



**GARTNER DISCLAIMER**

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.