

# Preempt

*Autonomous AI-Led Offensive Security and Validation*

## Offense-Led Security Operations

Attackers automate reconnaissance and weaponize disclosures in days. Defenders still run pen tests once a quarter. The window between a new exposure and active exploitation has collapsed and quarterly reviews cannot close it.

Atlas Preempt resolves this with an offensive AI operative that runs adversarial exposure validation, correlates against vulnerability signal from any scanner in your stack, and maps exposures to the attack paths most likely to be used against you. Every offensive action runs inside engineered human-judgment controls.

### -7 Days

Mean time to exploit

*Mandiant M-Trends, 2026*

### 53%

Of security leaders cite qualified candidates as a high-impact challenge

*KPMG Cybersecurity Survey, 2026*

### 10-20%

Utilization of existing security technology

*Ernst & Young, 2025*

## Continuous Offensive Point-of-View

Pen tests give you a snapshot. Atlas Preempt gives you a feed. The platform runs attack simulations and validation against your environment on a recurring cadence to help prioritize what an attacker would reach for next. Offense sharpens detection. Detection accelerates response. The loop tightens with every cycle.



## ATLAS PREEMPT DELIVERS

### Hundreds

Recon, vuln, exploit, and agentic hacking tools orchestrated by Atlas

### 5-Stage

CTEM lifecycle covered end-to-end on a single platform

### Weeks → Hours

From vulnerability disclosure to validated risk in your environment

### 100%

Reversible, policy-bounded offensive actions, every one logged

## 5-STAGE CTEM LIFECYCLE

One platform across all five stages - no tool stitching, no orphan findings. Evidence packaged for boards, regulators, and cyber insurers.

Scoping



Discovery



Prioritization



Validation



Mobilization



## AI GUARDRAILS

### Trust, Built-In to Preempt

Atlas Preempt runs on an offensive AI operative governed by the same trust conditions that hold across the entire Controlled Autonomy SecOps platform. Atlas does not act outside the customer's defined authority envelope.

- 1 Explainability**  
Every simulated attack produces a human-readable rationale, what was tried, why, and what was learned.
- 2 Shadow Approval**  
High-consequence actions that are destructive or outside test scope, are staged for human review with the investigation already complete.
- 3 Policy-Bounded Authority**  
The customer defines what Atlas can probe, when, and how aggressively. Atlas executes inside that envelope.

## HOW IT WORKS

Preempt is not a standalone audit. It is the offense half of a closed security lifecycle, findings from Preempt sharpen the detection content, validate exposures for reprioritization, and downstream remediation activities.

The screenshot displays the Atlas Preempt interface for an exploit titled "Exploit: PT-2026-001". The interface is dark-themed and includes several panels:

- Summary:** A critical SQL injection vulnerability was identified in the login form at portal.acme.com/auth/login. The username parameter is not properly sanitized, allowing an attacker to inject arbitrary SQL commands.
- Details:** Severity: CRITICAL; CVSS Score: 9.8; CWE: CWE-89: SQL Injection; OWASP: A03:2021 - Injection; Protocol: HTTPS; Port: 443; Associated Threats: SQL Injection, OWASP Top 10, Authentication Bypass, Data Exfiltration; Detection: D-260401-A1B2.
- Proof of Concept:** Shows steps for reconnaissance and probing for an injection vector. The probe step shows a request with a single quote in the username field, resulting in a MySQL syntax error.
- Detection Coverage:** Shows "Exploit Detected" with a 4 min time to detect, detection ID D-260401-A1B2, and detected at 2026-03-10 14:26 UTC.
- Remediation:** Priority: Immediate. Steps include replacing string concatenation with parameterized statements, applying input validation whitelists, enabling WAF SQL injection rules, and scheduling a retest.
- History:** A table showing the detection event on 2026-03-10 at 14:22 UTC, finding a critical issue.

## Why 2,000+ Security Teams Run with eSentire

*“We look at eSentire to be the experts. We trust them implicitly. They’re with us through the thick and thin till the end.”*

**CISO**, Financial Services  
GARTNER PEER INSIGHTS

*“eSentire has been an indispensable partner. Their proactive monitoring and strategic advice have contributed to a significant improvement in our cyber resilience.”*

**Security Leader**, Six-year customer  
GARTNER PEER INSIGHTS

2,000+ organizations in 80+ countries | 25 years of SOC expertise | 2M+ endpoints protected | MDR Leader, Forrester Wave™

## Who controls the pace, you or the attacker?

See how Atlas validates faster than the attackers can weaponize.

**LET'S TALK SECURITY →**

**IF YOU'RE EXPERIENCING A SECURITY INCIDENT OR BREACH CONTACT US  1-866-579-2200**

## eSENTIRE

eSentire is a leader in Controlled Autonomy SecOps, protecting 2,000+ organizations across 35+ industries around the world. Founded in 2001, the company's Controlled Autonomy SecOps operating model pairs agentic AI operatives with engineered human-judgment controls, delivering expert-depth security outcomes at machine speed without ceding accountability to opaque automation. Powered by the unified agentic AI Atlas Platform, eSentire's Atlas AI + 24/7 expert human SOC coverage delivers offensive capabilities that preempt exposures before attackers do, detect, and respond to stop threats in real time. For more information, visit [www.esentire.com](http://www.esentire.com) and follow [@eSentire](https://twitter.com/eSentire).