



PREVENT

AI-Assisted Attack Detection

Catches AI-assisted attacks while they're
being built, not after they land.

Every endpoint command, in real time

Verdicts, not alerts

Response before the payload runs

The detection gap


Attackers now move at machine speed. Most defenses still wait for something they've seen before.

Your organization runs more security tools than it did three years ago. Count the consoles, the dashboards, the alerts nobody reviewed because there were too many. Attackers count on that: they make noise to cover the real intrusion and walk between your defenses.


Now they have help. Commercial AI writes the reconnaissance scripts, the payloads, and the commands, then runs them faster than an analyst can open a ticket. The result is two clocks running at very different speeds.

EXHIBIT 1 Two clocks, drawn to the same scale

Breakout time is how long an intruder needs to move from the first compromised host to others. The breach lifecycle is how long an organization needs to identify and contain the breach.



The lifecycle rose to 247 days this year, **the first increase after five straight years of improvement.**¹ Anything that waits for a signature, a ticket, or an analyst is already late.



Industry average: 247 days to find and contain a breach.


PREVENT is built to catch the attack while it is still being assembled.

Al-assisted attacks in the wild

Confirmed campaigns and findings from the past twelve months. None of them are theoretical.

EXHIBIT 2 Twelve months of AI-assisted offense

October 2025 through September 2026. Each marker is a published campaign or industry finding.



JAN–FEB 2026

600+ firewalls, compromised by an amateur with AI

A financially motivated actor used several commercial AI services to compromise 600+ firewalls from a major OEM across 55+ countries in under six weeks. No vulnerability was exploited. Exposed management ports and weak passwords were enough, worked at a scale that once took a skilled team.

Amazon Threat Intelligence³

FEBRUARY 2026

AI-generated scripts on the endpoint

Attacks by AI-enabled adversaries rose 89% in a year. One eCrime group used AI-generated scripts to speed up credential dumping and erase forensic evidence. Average breakout time fell to 29 minutes.

CrowdStrike Global Threat Report²

AUGUST 2026

Payloads and shell commands, written by AI

Threat actors used AI to generate payloads and shell commands. One campaign sent nearly 200,000 AI model requests in two minutes. Some state-linked groups exploited critical vulnerabilities within 24 hours of a public proof of concept.

CrowdStrike Threat Hunting Report⁴

THE COMMON THREAD

Different actors, different targets. But each campaign ends the same way: **commands a machine wrote, running on an endpoint.** That is where PREVENT watches.

WRITTEN
BY A PERSON


BUILT
BY A MACHINE

THE SOLUTION

How PREVENT sees an AI-assisted attack

It judges how a command was built, not whether it has been seen before.

Endpoint tools built on signatures and known-bad behaviors are reactive by design: they recognize what they have already met. PREVENT's AI Command Detector asks a different question. It identifies when an attacker, or an AI working for one, is building, staging, or executing an attack in real time, even when the specific technique has never been seen before.



Unseen techniques

Built to catch attacks no signature has met yet.

Mapped to MITRE ATT&CK

Verdicts tie to the techniques your team already tracks.

Native to the agent



Runs inside PREVENT on every endpoint. Nothing extra to deploy.

The other side of AI security

The industry is busy securing its own AI. PREVENT watches for the attacker's.

EXHIBIT 3 Whose AI, and where it shows up

Securing AI models, governing LLM use, and finding shadow AI is important work. It covers one column. Attackers are using AI to generate payloads, obfuscate commands, and automate exploitation at machine speed, and that column has far fewer watchers.



Verdicts, not alerts

More alerts are not the answer to faster attackers. PREVENT's detections reach your team as decisions, not questions. Each one arrives with a verdict and the response already taken, so analysts review outcomes instead of working a triage queue.

The detector is native to the PREVENT endpoint agent. There is no separate model to deploy and no retraining cycle to wait on.

- No quarterly retraining
- No bolt-on model
- Auto-remediation through Commander


Anatomy of an AI-assisted attack

The same LLM-built command meets two kinds of defense. Only one of them stops it.

05

EXHIBIT 4 One attack, two outcomes

The attack moves left to right. Each lane shows where a defense gets its chance to act.



THE ATTACK

An attacker uses an LLM to generate an obfuscated PowerShell payload. The command is syntactically perfect: no typos, no hesitation, no trial and error.

It chains encoded instructions, manipulates the environment, and targets a system directory. **A signature-based tool sees a PowerShell command, finds no match, and lets it run.**

PREVENT'S RESPONSE

PREVENT recognizes the marks of machine-built tradecraft: the kind a human operator doesn't leave and a signature can't see. **Verdict: malicious.**

Commander isolates the endpoint, kills the process, and pushes a fleet-wide policy update, **before the payload executes.**

SIGNATURE-BASED TOOLS ASK


"Have I seen this before?"

PREVENT ASKS

"Was this built by a machine?"

PREVENT in the Beacon platform

Four capabilities built to work as one loop, from the open internet to the endpoint and back.



CYFAX tells you the door is open.

PREVENT tells you someone walked in.

ARETE tells you who, why, and who's next.

ARETE, Beacon's predictive analytics capability, reads signals from the criminal underground as early indicators of a campaign in preparation. The research behind it is published: *Underground Economy Signals as Temporal Predictors of Cybersecurity Breaches*, SSRN, doi:10.2139/ssrn.6584698.

CYFAX

Observes exposure outside your walls across 20,000+ criminal sources and 9,000+ tracked threat actors.

PREVENT

Recognizes the attack on the endpoint, validates it, and stops it.

ARETE

Forecasts who is building a campaign against you, and when the window opens.

Commander

Acts across the fleet: isolates, kills, and pushes policy.

Included, not added on. AI-Assisted Attack Detection is part of every PREVENT deployment, for current and new customers, at no additional cost.

Schedule a demo

See AI-Assisted Attack Detection working against real-world attack patterns in a 30-minute session with our engineering team.

Request a CYFAX assessment

See your external exposure the way an attacker does: leaked credentials, exposed services, and risk, in a branded report. No agent required.

detect.solutions/beacon-prevent inquiries@gobeacon.ai

SOURCES

1. IBM and Ponemon Institute, *Cost of a Data Breach Report 2026*, July 2026.
2. CrowdStrike, *2026 Global Threat Report*, February 24, 2026.
3. Amazon Threat Intelligence, AWS Security Blog, February 2026.
4. CrowdStrike, *2026 Threat Hunting Report*, August 2026.



BEACON TECHNOLOGY GROUP

CYFAX | PREVENT | ARETE

Beacon Technology Group | GreyIP Technologies Inc.

Miami, United States

USPTO Patent: AI-Based Threat Processing

detect.solutions/beacon-prevent

inquiries@gobeacon.ai

Rev. September 2026

© 2026 GreyIP Technologies Inc. All rights reserved.