INTRODUCING

AI Assisted Attack Detection

The first endpoint security system that detects AI based attacks while they're being built.

14+	3	15	0
Behavioral Signals	Detection Tiers	MITRE ATT&CK Techniques	False Positives

PREVENT by Beacon Technology Group

We don't detect attacks after they happen.
We detect them while they're being built.

BEACON TECHNOLOGY

April 2026 | detect.solutions/prevent

01

THE PROBLEM

Reactive Cybersecurity Is Dead.

Your organization has more security tools than it did three years ago. Is it more secure?

Count the active consoles. The dashboards. The alerts no one reviewed because there were too many.

Attackers know this. They launch noisy attacks to camouflage the real infiltration while your analysts are drowning in noise. They don't break through your defenses — they walk between them.

If your tools don't talk to each other, your team works for the software instead of protecting the organization.

76

Average security tools per enterprise

Panaseer 2024

62%

Of alerts go uninvestigated daily

Trellix 2024

277

Average days to identify a breach

IBM Cost of Breach

Industry says 277 days to identify a breach.

ARETE says 6–21 weeks *before* it happens.

Organizations using extensive security AI and automation saved roughly $1.9 million in breach costs compared to those that did not.

— IBM / EY, February 2026

> *"Fragmented cybersecurity is failed cybersecurity. Attackers only need one blind spot. You need the complete picture."*

That is why we built something that changes how a security operations center works.

Fewer tools. More intelligence. Less noise. More direction.

THE THREAT LANDSCAPE

AI-Assisted Attacks Are Here. Now.

These are not theoretical risks. These are confirmed incidents from the past twelve months.

FEB 2026 | AWS REPORT

600+ Firewalls Breached by AI-Armed Amateur

A Russian-speaking threat actor with limited technical skill used commercial LLMs to breach 600+ next-gen firewall instances across 55 countries in five weeks. AI automated reconnaissance, credential harvesting, and lateral movement planning. The actor targeted backup infrastructure for ransomware staging.

SEP 2025 | FIRST AI-ORCHESTRATED CAMPAIGN

AI Agent Executes 80–90% of Espionage Campaign

A state-sponsored group manipulated an AI coding agent to infiltrate roughly thirty global targets. The AI performed reconnaissance, credential harvesting, lateral movement, and data exfiltration with only 4–6 human decision points per campaign. Attack speed: thousands of requests per second.

Anthropic Threat Report, 2025

FEB 2026 | CROWDSTRIKE GLOBAL THREAT REPORT

89% Surge in AI-Enabled Attacks

Average breakout time dropped to 29 minutes in 2025, with some attacks completing in seconds. AI tools are being used across the full kill chain — reconnaissance, exploitation, evasion, and persistence. CrowdStrike calls it the defining shift of the current threat landscape.

FEB 2026 | IBM X-FORCE THREAT INDEX

44% Increase in AI-Enabled Exploitation

Attacks beginning with exploitation of public-facing applications surged 44%, driven by AI-enabled vulnerability discovery. Active ransomware groups increased 49% year over year. IBM warns that multimodal AI models will enable adversaries to automate complex tasks like advanced ransomware attacks.

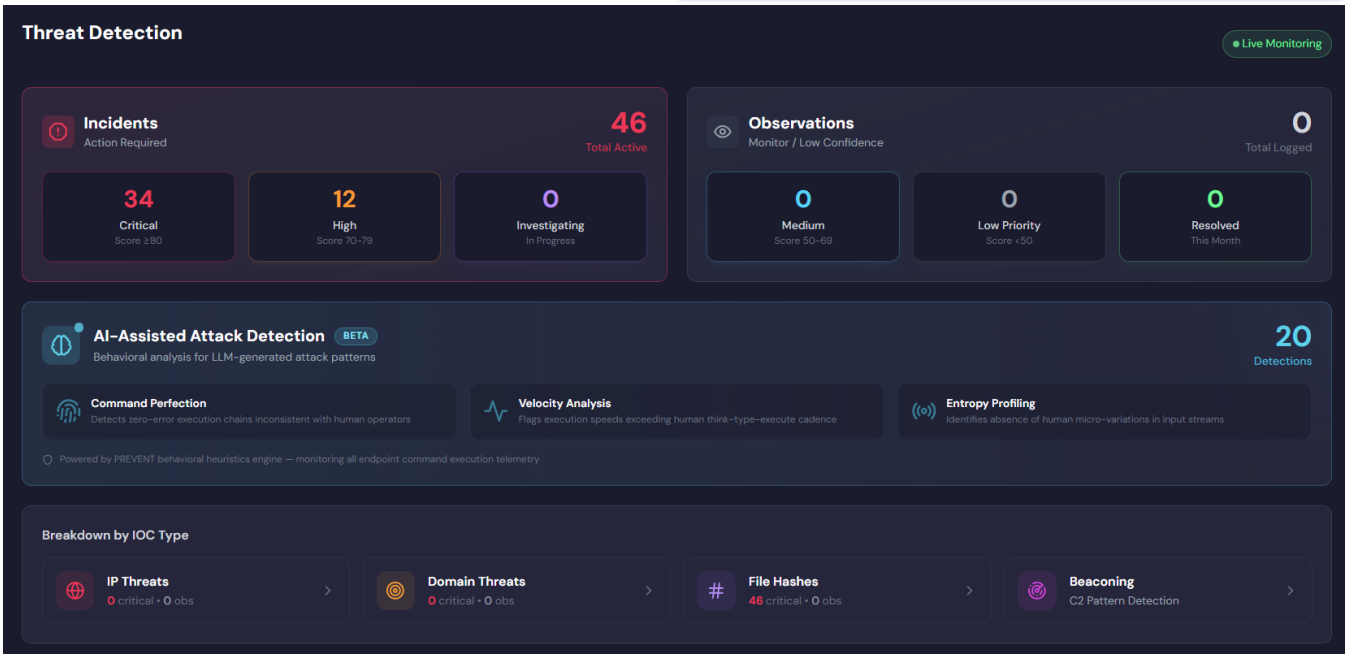

PREVENT Threat Detection Dashboard — AI-Assisted Attack Detection with 20 live detections in production

87% of organizations experienced an AI-driven cyberattack in the past year. **82.6%** of phishing emails now use AI in some form.

Cisco 2025 Cybersecurity Readiness Index | Keepnet / VIPRE 2025

02

THE SOLUTION

AI Assisted Attack Detection

Current endpoint tools detect known malware signatures and known-bad behaviors. They are reactive by design. Our AI Command Detector identifies when an attacker — or an AI working for an attacker — is building, staging, or executing an attack in real time, even if the specific technique has never been seen before.

THREE-TIER DETECTION ARCHITECTURE

TIER 1 — COMMAND ANALYSIS

Every command executed on the endpoint is decomposed into structural components: encoded payloads, obfuscation patterns, pipe chains, environment variable manipulation, registry modifications, and 14+ behavioral signals — each scored independently.

TIER 2 — COMPOUND SCORING

Individual signals are weighted and combined into a compound threat score. A single obfuscated command might score low. An obfuscated command with encoded payload, piped to a download cradle, targeting a system directory — that compound pattern triggers escalation.

TIER 3 — CONTEXTUAL VERDICTING

The compound score is evaluated against 47 monitored RMM processes, known-good baselines, and MITRE ATT&CK technique mappings across 15 categories. The verdict — benign, suspicious, or malicious — is rendered in real time with zero false positives in production.

WHY THIS ISN'T ANOTHER BOLT-ON

What everyone else is doing

The industry is focused on securing AI models — preventing prompt injection, detecting shadow AI usage, governing LLM deployments. Important work. But it ignores the other side of the equation entirely:

Attackers are using AI to generate attack payloads, obfuscate commands, and automate exploitation at machine speed. No one is watching for that.

What Beacon built

PREVENT's AI Detector is native to the endpoint agent. It analyzes every command in real time against 14+ behavioral signals, baselines 47 RMM processes, and renders compound verdicts with zero false positives.

No quarterly retraining. No analyst triage. No bolt-on model. Auto-remediation via Commander.

03

THE DIFFERENTIATOR

ARETE AI Predictive Analytics

Not a risk calculator. A threat actor campaign simulator that runs in reverse.

JULY 2025 — RESEARCH PAPER

- 27 major incidents analyzed retrospectively (Jun 2024 – Jun 2025)
- 86% validation rate across 6–21 week predictive windows
- 90–99% prevention probability across most cases
- Known gap filed: zero-credential attacks (e.g., Split Airport)
- Known gap: private IAB-TA channels outside marketplace visibility

MARCH 2026 — PRODUCTION

- 4 ransomware attacks predicted prospectively — not back-tested
- 4 for 4 validated by ransomware.live within 72 hours
- Zero-credential gap CLOSED — IKRON: 98% confidence, 0 credentials, Lockbit5 confirmed
- Marketplace provenance + threat actor tooling = 3D scoring
- **Model is outperforming the research paper, LIVE**

THREE LAYERS UNDERNEATH THE SCORE

Signal Detection

What's exposed — credentials, exploitable services, email weaknesses, typosquat domains, web app vulnerabilities. This is what the dashboard shows.

Marketplace Provenance

Where the signal transacted and which threat actor groups historically purchase from those channels. Invisible to the end user. Drives the probability score.

Operational Tempo Matching

Known scaffolding time from acquisition to detonation for specific groups. Qilin exploits AnyDesk + Splashtop. Lockbit5 targets exposed RDP. ARETE matches infrastructure to weapon.

ARETE doesn't ask "how exposed are you?" — it asks "is someone building a campaign against you, and how far along are they?"

92%	6–21	3,600+	500B+
Efficacy Back-tested	Weeks Predictive Lead Time	Threat Actors Tracked	Objects Ingested

THE PLATFORM

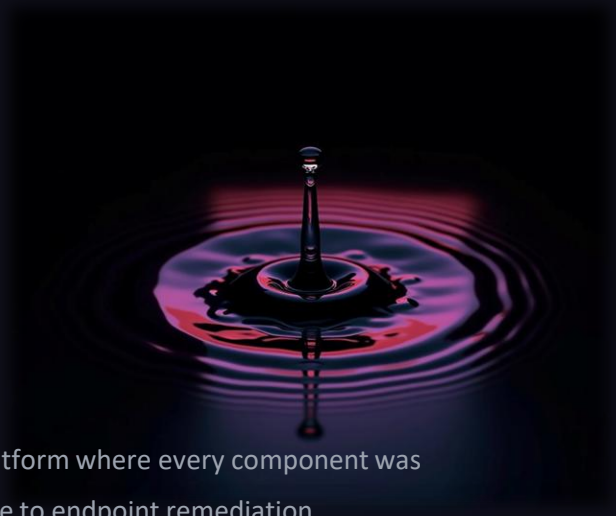

Fewer Tools. More Intelligence. Less Noise. More Direction.

PREVENT is not a collection of features bolted together. It is a single platform where every component was purpose-built to work as one system — from external threat intelligence to endpoint remediation.

ONE PLATFORM REPLACES

Vulnerability Scanner	→	**60 CIS/MITRE-mapped security controls**
EDR / Endpoint Agent	→	**Lightweight agent with AI Command Detector**
RMM Tool	→	**Commander fleet execution + auto-remediation**
NDR / Network Monitor	→	**Native NDR via ETW integration**
Dark Web Monitor	→	**CYFAX — 20,000+ criminal sources**
Threat Intel Platform	→	**ARETE — predictive analytics, 3,600+ actors**
EASM / Attack Surface	→	**Continuous external posture scanning**

Closed-Loop Architecture: CYFAX discovers threats externally → PREVENT validates and remediates on the endpoint → ARETE predicts what comes next → Commander auto-remediates across the fleet.

20+	15+	>60%	2
Capabilities Integrated	Point Solutions Replaced	Cost Reduction	Capabilities With No Equivalent

HOW IT WORKS

Stopping an Attack in Progress

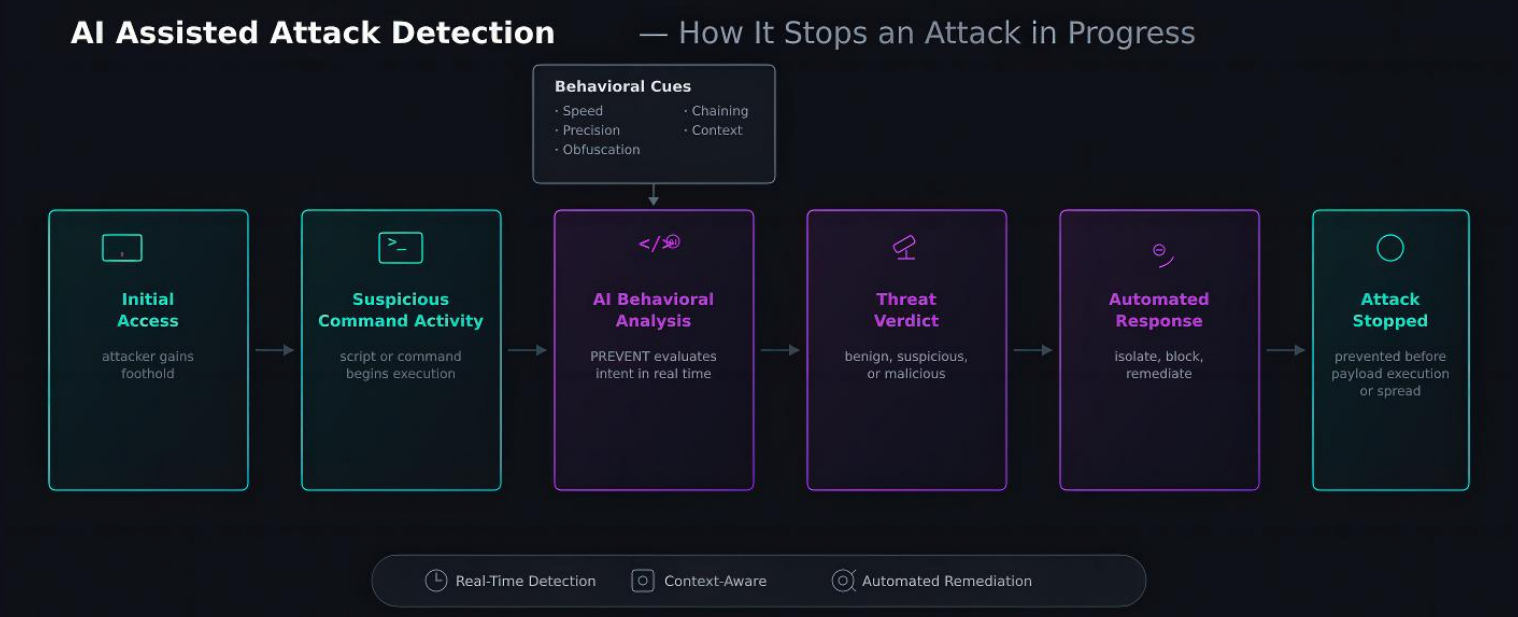

THE ATTACK

An attacker uses an LLM to generate an obfuscated PowerShell payload. The command is syntactically perfect — no typos, no hesitation, no iterative debugging.

It chains encoded instructions through pipe operators, manipulates environment variables, and targets a system directory.

Traditional EDR sees a PowerShell command. Checks signatures. Nothing matches. The command executes.

This is how 62% of alerts go uninvestigated. The payload looks clean to signature-based tools.

PREVENT'S RESPONSE

PREVENT watches for five things no human attacker does — and no signature catches:

- Commands executed with zero errors, zero hesitation
- Execution speed faster than any human can type
- Input patterns with no human micro-variation
- Payload obfuscation layered beyond normal admin work
- Privilege escalation disguised as routine operations

Compound score exceeds threshold. Verdict: Malicious.

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

Traditional EDR asks: "Have I seen this before?" PREVENT asks: "Was this built by a machine?"

CYFAX tells you the door is open.
PREVENT tells you someone walked in.

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

No other vendor in the world can say that.

Available Now

AI Assisted Attack Detection ships with PREVENT v2.1 — included for all current and new customers at no additional cost.

The capability that no other endpoint vendor offers is now part of every PREVENT deployment.

Schedule a Demo

See AI Assisted Attack Detection running live against real-world attack patterns. 30-minute session with our engineering team.

Request a CYFAX Assessment

Run your domain through CYFAX in 30 seconds. Branded report with leaked credentials, risk score, and predictive timeline. No agent required.

https://detect.solutions/beacon-prevent | sales@cyfax.ai

Beacon Technology Group | GreyIP Technologies Inc.
Miami, Florida | Veteran-Owned | USPTO Patent: AI-Based Threat Processing

April 2026