CHEAT SHEET

AUTONOMOUS EXPOSURE REMEDIATION · 2026

The no-need-to-cheat sheet for autonomous remediation.

Taking a closer look at Gartner's 2026 Hype Cycle for Security Operations.

7	5	1
Questions to vet a vendor	Gartner cautions, mapped	Verdict: real, or a demo

SOURCE
Gartner Hype Cycle for Security Operations · 2026

2026
Autonomous Exposure Remediation

WHERE WE ARE

AER just landed on Gartner's Hype Cycle

Gartner placed Autonomous Exposure Remediation at the Innovation Trigger in 2026 — high benefit, early maturity. That gap is exactly why diligence matters.

Illustrative recreation · Gartner Hype Cycle for Security Operations, 2026.

WHAT THIS MEANS

Every vendor will now claim it. Most are selling a demo.

Under 1% adoption, five to ten years to plateau — but a high benefit rating. This sheet is how you separate real autonomous remediation from a demo, before you hand software permission to change production.

Gartner
NAMED VENDOR

Gartner recognizes Reclaim Security as a sample Vendor across several related categories: Automated Security Control Assessment (ASCA), Dynamic Attack Surface Reduction (DASR), Preemptive Exposure Management (PEM), and Unified Exposure Management (UEM). The ASCA guidance maps directly to what we deliver: build on the incumbents you already run, integrate endpoint, identity, and email, move past discovery, and pre-approve low-risk fixes.

Gartner does not endorse vendors; Sample Vendors are representative, not exhaustive.

THE DILIGENCE SHEET

The seven questions to ask any AER vendor

Print it. Bring it to the demo. A vendor who answers these cleanly is worth your time — one who dodges them is asking you to bet the business on a black box.

1 “What runs in production today — and which named customers measured the result?”

WHY IT MATTERS

An embryonic category is full of roadmap slides dressed as products.

RED FLAG

Logos with no outcomes, design partners under NDA, “we're in private beta.”

GREEN LIGHT

Named customers, measured results, and a timeframe.

HOW RECLAIM ANSWERS

Named customers, measured results, public case studies — not logos under NDA. All four below are published.

PROOF · ALL PUBLISHED

Competitive Power Ventures

70%

Higher resilience, in weeks

Telit Cinterion

80%

Ransomware scenarios cut · 90% less manual effort

PINEGATE RENEWABLES

47–64%

Ransomware & APT resilience, 4 months

2 “Is the LLM in the execution path — or only in the explanation?”

WHY IT MATTERS

Gartner says to benchmark LLM output and never assume it beats your current practice. A model that writes and runs production changes on the fly is the riskiest design in this category.

RED FLAG

“Our AI generates and applies the fix.” No separation between reasoning and action.

GREEN LIGHT

The model helps a human understand and decide; the change that executes is deterministic and repeatable.

HOW RECLAIM ANSWERS

The LLM navigates, explains, and reports. The remediation that runs against production stays deterministic and policy-bound: you approve the guidelines, and the system acts within them automatically, every time.

EXECUTION & VALIDATION

Does it prove it works — before it acts?

Validation and impact-prediction separate real remediation from a fix list.

3 “Before it acts, does it validate exploitability against my environment — or act on scanner output?”

WHY IT MATTERS

Acting on raw findings is how you get a flood of false positives and wasted change windows.

RED FLAG

The fix list is just the scanner list.

GREEN LIGHT

It proves an exposure is exploitable against your actual controls and business context before proposing a fix.

HOW RECLAIM ANSWERS

Asset-level exploitability validation against your specific controls and runtime — then a fix built for your environment, not a generic ticket pointing at a CVE database.

4 “How do you predict business impact before a change runs — and can you put a number on it?”

WHY IT MATTERS

The fear that blocks remediation isn’t the exposure. It’s breaking the business with the fix.

RED FLAG

“We test in staging” or “we roll back if something breaks.” That’s reaction, not prediction.

GREEN LIGHT

A model that forecasts which users, workflows, and dependencies a change touches — before it runs.

HOW RECLAIM ANSWERS

PIPE™ — the Productivity Impact Prediction Engine — predicts business impact with 99.7% accuracy before any change executes.

ADOPTION & OUTCOMES

Can you adopt it safely — and does it close the loop?

Trust is earned in stages, and discovery only matters if the fix actually ships.

5 “Can I run it in audit mode and dry-run against my data first — and what are the rollout stages?”

WHY IT MATTERS

Gartner warns against expecting instant results and tells you to plan a multiyear adoption. Trust in automation is earned in stages, not granted on install.

RED FLAG

The only modes are off and autonomous.

GREEN LIGHT

A clear ladder — audit, dry run, small ring with human approval, widening scope, then autonomy for the change classes you choose.

HOW RECLAIM ANSWERS

Audit mode, dry run against production data, ring-based rollout, human-in-the-loop, then autonomous for the classes you decide to trust. You earn autonomy one validated step at a time.

6 “If you automate discovery, do I just get a bigger backlog — or do you close the loop?”

WHY IT MATTERS

Gartner’s sharpest caution: automating discovery without fixing prioritization and remediation builds a chokepoint — a faster firehose into the same clogged drain.

RED FLAG

The product finds more, faster — and still hands you tickets.

GREEN LIGHT

The product executes the fix across control planes, not just routes work to a queue.

HOW RECLAIM ANSWERS

Reclaim is the execution layer, not another scanner. When an attack chains across mail, identity, and endpoint, it composes one governed remediation plan across those planes — and predicts the impact of the whole chain.

FIT & HONESTY

Does it fit your stack — and tell you the truth?

The last question on cost, and the bonus that separates honest vendors from hype.

7 “Does this run on my existing stack — and what is the real cost?”

WHY IT MATTERS

Gartner tells CISOs to cost security into projects early rather than bolt it on.

RED FLAG

Rip-and-replace, or a tool that ignores what you already bought.

GREEN LIGHT

Runs on top of your current stack and makes those tools more useful.

HOW RECLAIM ANSWERS

Reclaim integrates with 40+ security platforms and turns them into inputs to one remediation motion. No rip-and-replace — the cost conversation is about getting more from tools already on the books.

RUNS ON TOP OF YOUR STACK

Defender · CrowdStrike · Entra · Microsoft 365 · Okta · Google Workspace · + 40 more platforms

THE BONUS QUESTION

What can your product NOT do yet? Where is the honest edge?

A vendor who claims the whole category is lying to you. AER as Gartner defines it reaches into code, software supply chain, and infrastructure-as-code inside agentic dev loops — a frontier ahead of nearly everyone today.

HOW RECLAIM ANSWERS

Our strength sits across the enterprise security stack: we surface vulnerabilities, misconfigurations, and control-plane exposures spanning endpoint, identity, email, and cloud configuration, where misconfiguration drives a large and growing share of breaches. We expand the same way we ask customers to adopt automation: prove it, then widen it.

NON-NEGOTIABLES

The governance checklist

Before any tool touches production, confirm it has:

- ✓ Approval gates you control, mapped to your change-advisory process, with least-privilege and just-in-time (PIM-based) access governing every remediation flow.
- ✓ A full audit trail of what changed, who approved it, and why.
- ✓ Rollback for every change.
- ✓ Read-only and guided modes — not just autonomous.
- ✓ A verifiable security and compliance posture (Reclaim is SOC 2).

THE MAP

How the questions map to Gartner's CISO cautions

Gartner caution to CISOs	Tested by	What good looks like
Maturity is low; plan a multiyear roadmap, not instant MTTR	Q1 · Q5	Named production customers plus a staged trust ladder
Expect false positives and low-quality fixes	Q3 · Q4	Exploitability validation and business-impact prediction before action
Benchmark LLM output; don't assume it's better	Q2	LLM out of the execution path; deterministic, repeatable changes
Don't build a remediation chokepoint	Q6	An execution layer that closes the loop — not another scanner
Cost security in early	Q7	Runs on the stack you own, resulting in provable ROI and lower TCO

THE ONE-LINE TEST

Demand to fix what others only flag.

If a vendor cannot keep the model out of the execution path, predict business impact before acting, and let you start in audit mode — you are not looking at autonomous exposure remediation.

You are looking at a demo.

- A named customer with a measured result
- The model kept out of the execution path
- Business impact predicted before acting
- A start in audit mode, autonomy earned

Companion to the Reclaim Security blog post on Gartner's Autonomous Exposure Remediation profile.

See the seven answers live.

Reclaim is the execution layer for exposure remediation — it sits on top of the 40+ security and identity tools you already run, validates what is actually exploitable, and executes governed fixes. Every change is simulated by PIPE™ (Productivity Impact Prediction Engine) before deployment, so nothing breaks the business.

RUNS ON TOP OF YOUR STACK

Defender · CrowdStrike · Entra · Microsoft 365 · Okta · Google Workspace · and 40+ endpoints

70%	80%	85%	99.7%
CPV resilience, in weeks	Telit ransomware scenarios	Aqua phishing exposure	PIPE™ impact accuracy

GARTNER RECOGNITION

Also named a Vendor in Gartner® — Automated Security Control Assessment, Hype Cycle for Security Operations, 2026.

Regain control. Rebuild trust. Reclaim Security.

reclaim.security/demo

Talk to us →