/exposure action platform™

Stop managing exposures. Start fixing them.

The Seemplicity Exposure Action Platform™ proves which exposures are real, neutralizes the ones that are, and carries the fix through to verified closure. It works through your own tools and change processes.

Exposure management stops at a prioritized, assigned backlog, and then a person still has to do the work. Seemplicity picks up where that stops. **AI Response Analysts, powered by AI Exposure Response Skills,** investigate each exposure, contain the real threats and drive permanent fixes to auditable closure.

> “The difference between an agent that guesses and an analyst you can trust is skill.”
>
> **Yoran Sirkis**
> Co-founder and CEO, Seemplicity

Action you can trust, grounded in context.

The Exposure Action Platform™ is made up of two modules that work together or stand alone.

AI Exposure Response (AER)

Is it real? Is it done?

AI Response Analysts prove whether an exposure is exploitable in your environment and contain it with the controls you already own. They then fix it through your change process and verify the result.

Runs on AEM or the exposure management platform you already use, including ServiceNow, Nucleus and Brinqa.

AI Exposure Management (AEM)

What matters? Who owns it?

AEM brings findings from 200+ scanners, cloud, code and CMDB tools into one clean backlog. Every finding is deduplicated, aggregated by fix, mapped to an owner, prioritized by risk and tracked to SLA.

Seema, the AI assistant, answers questions about findings, assets and SLAs in plain language.

AER does the fixing. AEM, or the platform you already run, tells it what matters.

Finding isn't the bottleneck. Fixing is.

AI is shrinking the time to exploit

AI now speeds up proof-of-concept exploits. As a result, "a public exploit exists" no longer separates the urgent few from the rest.

SLA backlogs can't keep pace

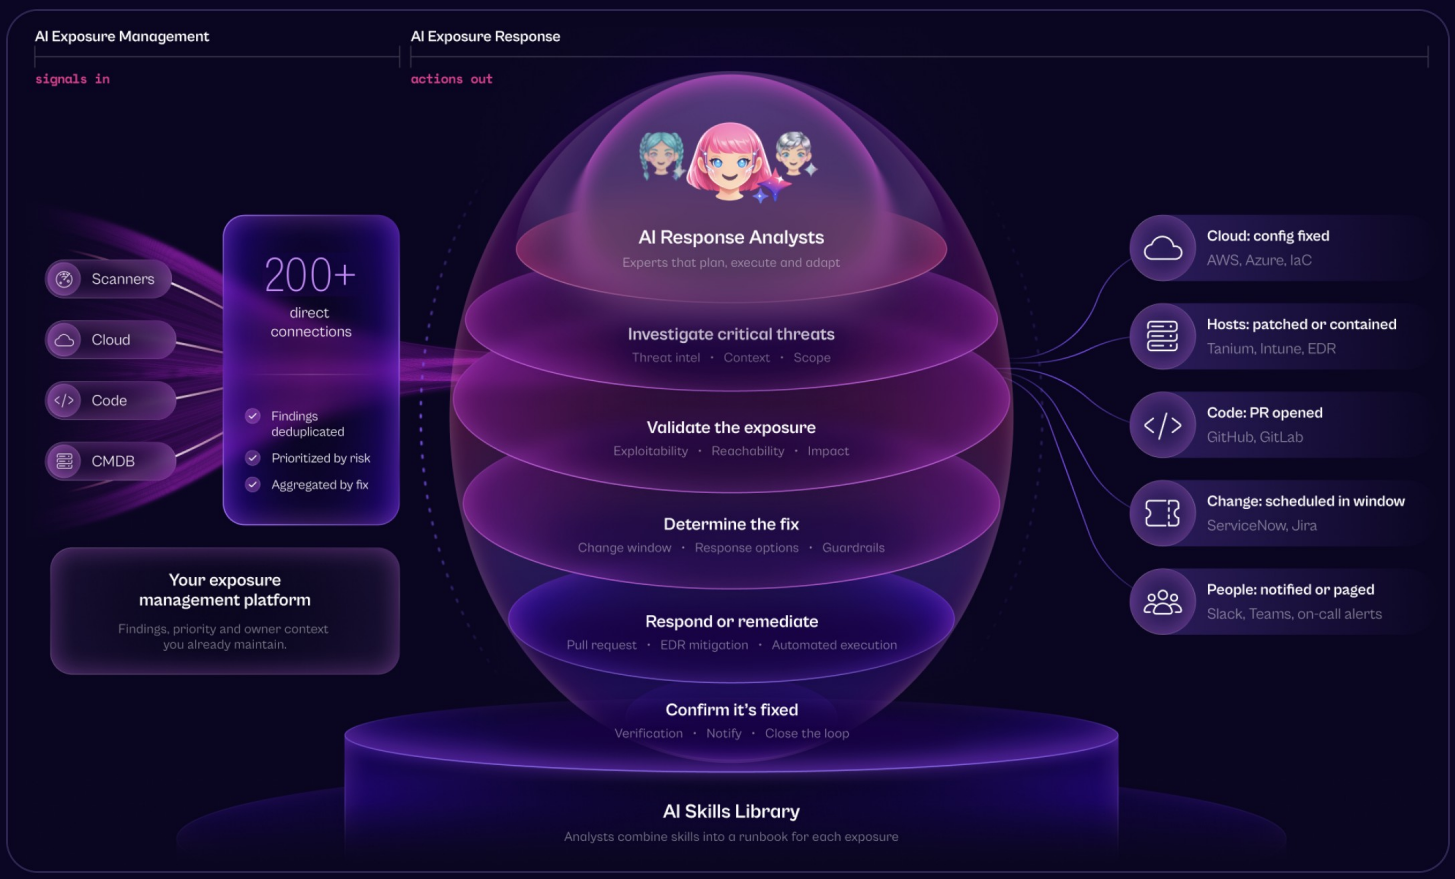
A better-ranked queue is still a queue. The fixing is still done by hand, one ticket at a time.

Generic agents can't act safely

A model and a loop lack the context, remediation expertise and safety judgment that production systems demand.

From found to fixed, on one platform

AI Exposure Response (AER) does the work, grounded in the findings, owners and priorities that AI Exposure Management (AEM) delivers.



01 Investigate

Proves whether it's real. Where it isn't, it's deprioritized with the evidence attached.

02 Contain

Where it is real, it's contained now with the EDR, WAF or firewall you own.

03 Fix

The permanent fix follows in your change window, through your own process.

04 Verify

Confirms it's resolved, with proof, and closes the loop.

Your best engineer, on every finding

A skill is a unit of security expertise an AI agent can carry out. It captures **the steps an experienced engineer takes to answer one question about an exposure**, the tools used to verify the answer, and the evidence needed before acting.

`/exploitability`

Do the prerequisites in public proof-of-concept exploits hold in your live environment?

`/reachability`

Is the vulnerable code actually loaded and called?

`/root cause`

Which repository, build or base image introduced the vulnerable component?

`/ownership`

Which team owns it, based on ticket history, git logs and assignment records?

`/response path`

What is the fastest safe response: mitigate with a control like a WAF rule, remove an exploit precondition, or patch?

`/operational risk`

Does the fix need care, such as a kernel update that requires a reboot?

Every skill returns a verdict with an expandable evidence trail, so each decision can be reviewed and audited.

Five analysts, ready on day one

Each analyst packages skills for one part of the attack surface. Each one returns a verdict, evidence, priority, fix options and an owner.



Host VM Analyst

`infrastructure vulnerabilities`

Checks exploit prerequisites against live hosts through EDR. It also tests network reachability and flags fixes that need a reboot.



SAST Analyst

`code vulnerabilities`

Reads source code in GitHub and GitLab to confirm whether SAST findings are reachable. It then sizes each fix, so developers only see validated, scoped work.



SCA Analyst

`dependency vulnerabilities`

Confirms that vulnerable library functions are actually called. On a zero-day, one SBOM-powered query finds every instance before scanner signatures catch up.



Web Analyst

`web vulnerabilities`

Maps DAST alerts and bug bounty reports to the code behind them, even behind a gateway. It names the owner and can open a draft fix PR.



Container Analyst

`container vulnerabilities`

Traces each vulnerable package to the build and layer that introduced it. It then recommends a fix for the root cause.

Your runbooks, as skills

`your team's expertise`

Write a runbook in plain language and it becomes an agent: every finding fixed your way, and verified.

Three layers a generic agent doesn't have

generic agent = a model and a loop

layer 03
Specialized harness
A governed runtime that knows a reboot needs a change window and every change needs a way back.

layer 02
Specialized skills
Remediation steps written by people who have done this work, across the whole tool stack.

layer 01 / foundation
Specialized data
Context for every finding: normalized across tools, mapped to owners, with live evidence.

Already running at customers

87%
of critical findings proven not exploitable

75%
of unowned bug bounty reports mapped to source code

See it on your own findings

AI Exposure Response Skills are available now in Seemplicity's AER and AEM modules. Buy either module or both. AER also runs on your existing exposure management platform, including ServiceNow, Nucleus and Brinqa.

Request a demo

seemplicity.ai/demo

/trusted by the world's leading enterprises

Carlsberg Group · Morningstar · SoFi · Graphic Packaging International · NTT · West Bend · Hertz · OpenWeb

© 2026 Seemplicity · 04 / 04