mondoo

FIND IT. FIX IT. PROVE IT.

Can your security program do all three?

A GUIDE FOR SECURITY AND PLATFORM TEAMS

Every security tool finds problems and hands you the *list*.

MONDOO FINDS THEM AND *FIXES* THEM.

Mondoo is an agentic security platform. AI agents work continuously across everything you run, find the vulnerabilities, misconfigurations and unapproved AI in your estate, ship the fix through the tools you already have, and prove them closed. Your team sets the thresholds and approves the changes. Mondoo can run the whole thing for you, or you can run it yourself on the platform. Most organizations start with one and move toward the other as their confidence grows.

WHY THIS MATTERS NOW

62% of teams still remediate manually, according to Mondoo's 2025 State of Vulnerability Remediation. That is why backlogs grow faster than anyone can work them, and why findings raised months ago are still open. Not because nobody noticed. Because nobody had capacity to close them.

THE LOOP

01
DETECT

Continuous, agentless assessment across cloud, on-premises, endpoints, SaaS, network devices, containers, the software supply chain, and the AI tools and agents your people installed themselves.

Agentless matters, because coverage that waits on a deployment project is coverage you do not have yet.

02
PRIORITIZE

Ranked by business impact and real exploitability rather than severity score alone. A critical CVE on an isolated system is not the same risk as a medium on something internet-facing holding customer data, and a queue that cannot tell the difference wastes the capacity you have.

03
SHIP

The fix itself. Code changes, pull requests, and configuration changes delivered through Intune, Jamf, Ansible, GitHub, Terraform and your ticketing system, with the context an engineer needs to act. Nothing reaches production without your approval, and every action is logged. This is the step that separates a program that reduces risk from one that reports it.

04
VERIFY

Confirmation that the fix landed and holds, captured as evidence. Without this stage you have a belief about your risk rather than a fact.

WHAT SITS UNDERNEATH

70+ integrations

Coverage that reaches both the systems you assess and the tools you remediate through.

783,000+ vulnerabilities and malicious packages

A vulnerability intelligence database spanning every major ecosystem, updated continuously.

AI agents and skills

Every AI agent, plugin, MCP server and skill on your fleet, inventoried from the endpoint and governed before it runs. Of the 58,000+ agent skills we have scanned, roughly one in five came back clean.

Policy as code

What compliant means in your organization is written down, versioned, and applied the same way every time rather than interpreted per audit.

The frameworks you answer to

Assessment against SOC 2, ISO 27001, NIS2, DORA, PCI DSS, HIPAA, CIS Benchmarks and DISA STIGs.

WHAT CHANGES

Across managed estates, closing the loop produces three measurable shifts. Not findings triaged. **Findings closed, proven, and held closed.**

60%

Fewer open vulnerabilities

Inside the first 90 days.

<16D

Mean time to remediation

On critical issues, where the same work done manually runs to months.

10X

Faster than doing it yourself

Without adding headcount to triage and patch.

WHAT IT LOOKS LIKE IN PRACTICE

Deutsche Telekom

More than 10,000 assets under unified visibility across a group estate.

> "The speed and accuracy of Mondoo's platform, combined with its deep insights into IT architecture, enables customers to quickly remediate issues and significantly reduce CVEs and policy violations."
>
> - **Thomas Tschersich**, CEO Telekom Security and CSO, Deutsche Telekom AG

European travel group

Around 1,000 assets, from AIX to AWS, under continuous assessment. Six people operate a program that would otherwise need ten. **Audit evidence that used to take months now takes three to four days.**

Campminder

Cut vulnerabilities by 50%, retired four separate tools, achieved PCI DSS, and **reduced manual work by four to five times,** with a team that stayed the same size.

THE PATTERN ACROSS ALL THREE IS THE SAME, AND IT IS WORTH STATING PLAINLY: THE NUMBER WENT DOWN AND THE TEAM DID NOT GET BIGGER.

This should mean less work, not more

If you run platform, infrastructure, or operations, a security program closing more findings sounds like more tickets for you. It should be the opposite. Fixes arrive through the tools you already have rather than in a separate console you have to check. Nothing reaches production without your approval. You set the thresholds, and automation increases only as your confidence in it does. Every action is logged, so the audit trail is a by-product rather than a project. The queue shrinks rather than grows.

Evidence as a by-product

More people ask for proof than did three years ago. Boards want to know whether risk is going down. Auditors want evidence against NIS2, DORA, ISO 27001, SOC 2 and PCI DSS. Insurers price on it. Enterprise customers send questionnaires that assume you have it. A program that verifies every fix and holds it closed produces that evidence continuously. A program that stops at the report produces it in a scramble, three weeks before a deadline, from people who have other work.

HOW TO TELL WHETHER YOUR PROGRAM CLOSES THE LOOP

Nine questions. Most organizations answer the first few immediately and stall somewhere in the middle. Where you stall is the useful part. The questions are deliberately phrased as **could you produce this**, not *do you do this*. A program either has the evidence or it does not.

- [ ] **01 · Coverage**
  Could you produce a current inventory of everything you run, including on-premises and anything not managed by your cloud provider? How confident are you that it is complete?

- [ ] **02 · AI**
  Could you name every AI agent, plugin, or MCP server installed across your fleet, and say which of them can reach internal systems using an employee's credentials?

- [ ] **03 · Prioritization**
  Could you explain, for the top ten items in your queue, why they are the top ten in terms of exposure rather than severity score?

- [ ] **04 · Ownership**
  For a finding raised today, could you name the person who closes it and the system they close it in?

- [ ] **05 · Time to closed**
  Do you measure the time from finding to verified closed, or only the time to ticket? What was the median last quarter?

- [ ] **06 · Backlog trend**
  Is your number of open findings higher or lower than it was six months ago? Could you show the trend?

- [ ] **07 · Verification**
  When a fix is applied, is it re-checked and confirmed closed by something other than the person who applied it?

- [ ] **08 · Evidence on demand**
  If an auditor asked this week for evidence of continuous compliance against your primary framework, how long would it take, and who would stop their other work to assemble it?

- [ ] **09 · Capacity**
  If your finding volume doubled next quarter, what in your program would break first?

READING YOUR ANSWERS

Stalled at 01 or 02

The gap is coverage, and everything downstream is being prioritized against an incomplete picture.

Stalled at 04 to 07

The common problem: strong detection, no reliable path to closed. This is where most programs sit.

Stalled at 08 or 09

Your program works but it does not scale, and the evidence burden is carried by people rather than by the system.

WHERE TO TAKE THIS

Most teams start with an assessment. Mondoo maps what is exposed across your environment and shows what closing the loop would look like in practice, before anything changes.

MONDOO.COM