



FIND IT. FIX IT. PROVE IT.

Can your security program do all three?

A GUIDE FOR SECURITY, IT AND PLATFORM TEAMS

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

MONDOO FINDS THEM AND *FIXES* THEM.

Mondoo is the AI-native exposure management platform with agentic remediation. Its agents find the vulnerabilities, misconfigurations and unapproved AI across your code, cloud, endpoints and external attack surface, deliver the fix through the tools you already run, and prove it closed. Autonomous where it's safe. Approval where it matters. Run it yourself on the platform, or have Mondoo run it for you as Service as Software. Either way, autonomy grows as your confidence does.

WHY THIS MATTERS NOW

Around 66,000 new CVEs are forecast for 2026, more than 180 a day, according to FIRST. AI is accelerating vulnerability detection, and attackers can turn a new flaw into an exploit in minutes. Meanwhile 62% of teams still remediate manually, according to Mondoo's 2025 State of Vulnerability Remediation. That's why backlogs grow faster than anyone can work them. You can't prioritize your way out. The bottleneck isn't finding. **It's fixing.**

THE LOOP

01
DETECT

Continuous assessment across code and the software supply chain, cloud, on-premises, endpoints, SaaS, network devices, containers, AI and your external attack surface, including the AI tools and agents your people installed themselves. No new agents to roll out, so coverage doesn't wait on a deployment project.

02
PRIORITIZE

Every exposure is ranked by Mondoo Risk Dimensions: attack surface, blast radius, exploitability and business priority, not severity score alone. A critical CVE on an isolated test box isn't the same risk as one on an internet-facing server holding customer data, so internet-facing exposures move to the front of the queue.

03
SHIP

Agents build the fix with an impact statement and a rollback plan, and deliver it through the tools you already run: Microsoft Intune and CrowdStrike Falcon for endpoints, GitHub pull requests, the IDE and AI coding agents for code, Ansible and Terraform for infrastructure, and your ticketing system. Autonomous where it's safe, approval where it matters, and every action is logged. This is the step that separates a program that reduces risk from one that reports it.

04
VERIFY

Mondoo rescans to confirm the fix landed and records it as evidence. Policy as code fixes it again if it drifts. Without this step you have a belief about your risk rather than a fact.

NEW IN OCTOBER 2026

Remediation through Microsoft Intune and CrowdStrike Falcon

Fixes deployed to Windows endpoints through the tools IT already trusts, then verified.

Code inspection

Finds vulnerable code, the dependencies your code actually uses and exposed secrets, with fixes in pull requests, the IDE and AI coding agents. Source code never leaves your environment.

External attack surface scanning

Finds your internet-facing hosts from a single domain and checks TLS, DNS, HTTP, email security and known vulnerabilities. Every failed check comes with the fix.

Enhanced and extended CIS Benchmark remediations

Fixes for more failed checks, including on systems already in production, with enhanced remediations tested end to end on real Windows and Linux systems, and kept hardened by policy as code.

MCP

Your own agents and assistants can talk to Mondoo: ask what's exposed, pull the evidence and trigger the fix.

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.

Specialist agents, readable at every step

A team of task-specific agents, not one black box. Each explains what it did in plain language your team can check before anything changes.

The frameworks you answer to

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

WHAT CHANGES

Closing the loop changes what you measure. Not findings triaged.
Findings fixed, proven and kept fixed.

5M+

Remediations performed

By Mondoo's agents and experts, to date.

300+

Customers worldwide

From enterprises to managed security service providers.

ON DEMAND

Evidence of every fix

What was found, fixed, approved and verified, without the scramble before an audit.

OUTCOMES, NOT FINDINGS

WHAT IT LOOKS LIKE IN PRACTICE

Deutsche Telekom

More than 350,000+ 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.

Your path to autonomous remediation

If you run IT, infrastructure or operations, a security program closing more findings sounds like more work for you. It should be the opposite. Fixes arrive through the tools you already run, with an impact statement and rollback, and you decide how much runs on its own.

AUTONOMOUS

Agents fix, deploy and verify on their own, and log every change. For laptops, desktops and non-production systems.

APPROVAL

Agents prepare the fix and its impact statement; a person approves before anything changes. For production servers and databases.

HUMAN-LED

Agents plan the fix; your team runs it and builds a track record before automating.

WE DON'T HAND AI THE KEYS. WE HAND YOU THE CONTROL. AS YOUR CONFIDENCE GROWS, YOU TURN THE DIAL FURTHER.

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, PCI DSS and CIS. Insurers price on it. Enterprise customers send questionnaires that assume you have it. Mondoo records what was found, what was fixed, who approved it and when it was verified, so the evidence builds as you fix. 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?

☐ 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?

☐ 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?

☐ 07 · Verification

When a fix is applied, is it re-checked and confirmed closed by something other than the person who applied it, and fixed again if it drifts?

☐ 09 · Capacity

If your finding volume doubled next quarter, what in your program would break first? Which fixes could run on their own today, and which would still need a person?

☐ 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?

☐ 04 · Ownership

For a finding raised today, could you name the person who closes it and the system they close it in?

☐ 06 · Backlog trend

Is your number of open findings higher or lower than it was six months ago? Could you show the trend?

☐ 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?

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's exposed across your environment and shows what it could fix today, and where you'd want approval first, before anything changes.

MONDOO.COM

FIX NOW. FIX FOREVER.