DEFECTDOJO

EBOOK · 2026 EDITION
FOR SECURITY LEADERS

AI IN VULNERABILITY MANAGEMENT

IMPLEMENTING AI ACROSS ENTERPRISE VULNERABILITY MANAGEMENT

How AI becomes a force multiplier for security teams stuck in manual reporting: where it earns its place, how to keep your data safe, what executive briefings should contain, and how to get started.

500+	50M+	10,000+	IDC
OFFICIAL INTEGRATIONS	DOWNLOADS	ORGANIZATIONS	MAJOR PLAYER IN ASPM

Contents

01 Executive summary PAGE 02

02 Key challenges across enterprise vulnerability management PAGE 03

03 Where AI earns its place PAGE 04

04 Executive briefings, insights and recommendations PAGE 05

05 From findings to fixes PAGE 06

06 Why DefectDojo PAGE 07

07 Getting started with AI PAGE 08

01

Executive summary

Enterprise security teams are stuck in a loop of scanning, identifying and remediating vulnerabilities. Despite a wealth of scanning tools, much of their time goes to manual reporting and metrics instead of active testing, and duplicates and false positives across disconnected tools keep data clarity low.

This eBook explores how AI acts as a force multiplier: automating mundane triage and reporting, adding the data context needed for sound business risk decisions, and generating executive briefings, insights and recommendations.

You will see how the Model Context Protocol (MCP) gives AI the findings it needs without flooding it with raw data, and how to adopt AI without exposing sensitive security data, so your team moves from tactical data mining to strategic security architecture.

KEY TAKEAWAY

AI is only as good as the data under it. Normalize and deduplicate first, then let AI take the reporting and triage work, with a person making every final call.

02

Key challenges across enterprise vulnerability management

Every security layer a team adds (SAST, DAST, SCA, cloud) adds a flood of data that was never normalized, from tools that do not know about each other. Engineers end up mining data instead of doing the work they were hired for.

69%
OF SECURITY TEAMS REPORT METRICS MANUALLY OR MOSTLY MANUALLY

57%
SPEND 25% TO 50% OF THEIR TIME ON VULNERABILITY MANAGEMENT OPERATIONS

68%
LEAVE CRITICAL VULNERABILITIES UNRESOLVED FOR MORE THAN 24 HOURS

1 Time-consuming manual reporting and metrics

WHAT HAPPENS

Engineers spend most of their week mining data, building reports and calculating metrics by hand.

WHAT IT COSTS

Mundane work drives burnout and keeps people from the assessments, testing and remediation they were hired to do.

2 Limited bandwidth for assessments

WHAT HAPPENS

Red tape and manual data entry limit how many security assessments a team can finish in a year.

WHAT IT COSTS

Coverage stays on the crown jewels while the rest of the estate goes untested.

3 Prioritization on siloed data

WHAT HAPPENS

Scanners use different formats and field names for the same vulnerability, and several scanners on one repository produce duplicates and false positives.

WHAT IT COSTS

Teams compare apples to oranges and cannot tell which findings pose the most risk. 37% of organizations cite missing context as their top prioritization challenge.

Sources: SANS 2025 survey; Help Net Security.

03

Where AI earns its place

The main barrier to AI adoption is the fear of exposing sensitive security findings to public models. That risk is manageable. The payoff comes when AI works on clean, normalized data.

1 Automating triage and reporting

WHAT AI SHOULD DO

AI handles repetitive triage questions: is this likely a false positive, is it exploitable here, which findings share a fix. Engineers move off manual reporting and onto higher-value work.

HOW DEFECTDOJO PRO DOES IT

Triage Engine rules can ask AI to review findings (Sensei license). Each suggestion waits in AI Proposals until a person accepts or rejects it.

2 Data context without wasted tokens

WHAT AI SHOULD DO

Flooding a large language model with raw scanner output wastes tokens and fills the context window. AI should receive only the relevant findings, already normalized and deduplicated.

HOW DEFECTDOJO PRO DOES IT

A built-in MCP server lets AI assistants read, and work in, your DefectDojo data with each user's own permissions, so they ask for what they need instead of everything.

3 Enrichment for a prioritized inventory

WHAT AI SHOULD DO

AI gives better answers when findings already carry exploit intelligence such as EPSS and KEV. A raw list becomes a ranked inventory the model can reason about.

HOW DEFECTDOJO PRO DOES IT

EPSS, KEV and exploit-maturity data are added to every finding with a CVE and refreshed every 24 hours, feeding a priority score and risk band for every finding.

4 Tool stack rationalization

WHAT AI SHOULD DO

AI can show which tools produce the most false positives or recurring noise, so leaders can move budget from tool maintenance to strategic defense.

HOW DEFECTDOJO PRO DOES IT

Tool Insights shows which scanners earn their keep and which are the noisiest.

5 Keeping data under your control

WHAT AI SHOULD DO

Sensitive findings should never become training data, and AI access should never exceed what the person using it is allowed to see.

HOW DEFECTDOJO PRO DOES IT

Sensei is never trained on your data. The MCP server works with each user's own permissions, and DefectDojo Pro runs cloud-hosted, self-hosted or air-gapped.

04

Executive briefings, insights and recommendations

AI can take the tedium out of report writing, turning live vulnerability data into briefings for the SOC, the CISO and the board in minutes.

What a good AI briefing contains

- **Executive summary.** A high-level overview fit for a board presentation.
- **Threat intelligence brief.** Relevant threats, found by comparing vulnerability patterns with benchmarks such as the OWASP Top 10.
- **Emerging threat vectors.** The specific high-priority areas of concern.
- **Security posture and gap assessment.** Where defenses stand today, based on live data, and a structured path to a more secure state.
- **Strategic recommendations.** Initiatives that address high-priority risk, with success metrics to measure their impact.
- **SOC briefings.** Team assignments, workload distribution and immediate action items.

How DefectDojo Pro produces them

- The Command Center opens with a short daily brief that links straight to the evidence.
- A Posture Score from 0 to 1000 that leaders can follow back to the findings behind it.
- Dashboard templates, including an executive brief, that can send a scheduled executive pack as a PDF.
- Through the MCP server, the AI assistant your team already uses can draw on OWASP Top 10 and CWE reference material, run ready-made prompts such as a security landscape report, and build and schedule branded reports.
- Triage Engine rules can generate a report after each import and send the link to Slack or email.

A PERSON STILL OWNS THE CONCLUSION

AI drafts the briefing. Your team decides what it means. Every number in a DefectDojo report links back to the findings behind it, so a reviewer can check the claim before it reaches the board.

05

From findings to fixes

The next step for AI in vulnerability management is taking work off the queue entirely. Sensei, available as an add-on to DefectDojo Pro, is built for that.

Fixes as pull requests

Sensei scans code repositories and cloud accounts and opens pull requests that fix what it finds. You set the autonomy: review every change, or let it run.

Threat modeling before code exists

Threat Modeling reads a design document or architecture diagram and lists likely threats, attack paths and security requirements, which can be tracked in DefectDojo as findings.

Safer configuration changes

Sensei Advisor suggests changes to deduplication and prioritization settings and shows how each would affect your own findings. Nothing changes until someone applies it, and changes can be undone.

Governance for the AI your developers use

AI Inventory lists the AI coding assistants, MCP servers, models and AI packages found in your repositories, and AI Governance Policies raise a finding when an unapproved one turns up.

WHY IT MATTERS

An AI that summarizes a backlog is useful. An AI that ships the fix is what moves the backlog toward zero. Either way, control stays with your team.

06

Why the world runs on DefectDojo

Open source since 2014, born inside a real application security program, built with the OWASP community for the people who refuse to accept the status quo. No platform in security has been deployed, inspected and stress-tested more.

500+
OFFICIAL INTEGRATIONS

50M+
DOWNLOADS

10,000+
ORGANIZATIONS · OWASP FLAGSHIP

IDC
MAJOR PLAYER IN ASPM

One queue for every scanner

500+ integrations by file, API or scheduled Connector, all in one normalized model. **Most tools dedupe inside one scanner. DefectDojo Pro dedupes across every tool.**

Risk ranking you can defend to an auditor

EPSS, CISA KEV, threat intelligence, reachability and business context in one priority score, with weights you set per Asset. **Severity is not risk.**

Sensei ships the fix

Available as an add-on to DefectDojo Pro, Sensei opens pull requests that fix code findings and cloud misconfigurations. **You set the autonomy. Control stays yours.**

Evidence, not spreadsheets

Findings mapped to nine bundled compliance frameworks, with **FedRAMP POA&M and OSCAL output, CMMC scoring and EU CRA Article 14 reporting** built in.

Pricing that never punishes coverage

Unlimited users and assets. **Cloud-hosted, self-hosted or air-gapped, and your data belongs to you.**

Prove it before you buy it

Community Edition is the free open source platform, built on the same data model and parsers as Pro. **Run it in your stack, then bring your data to Pro when the program scales.**

CAPABILITY	COMMUNITY EDITION · WHERE PROGRAMS START	DEFECTDOJO PRO · WHERE PROGRAMS SCALE
Ingestion	360+ built-in parsers, by file upload, API or CI pipeline	Every Community Edition parser, plus the Universal Parser and 140+ Upstream Connectors that pull findings on a schedule
Deduplication	Same-tool deduplication, set in configuration	Cross-tool deduplication chosen per tool in the UI, plus Root Causes that group findings one fix will clear
Prioritization	Ordered by severity and CVSS, with business attributes recorded on each Asset	A priority score and risk band for every finding, weighing EPSS, KEV, threat intelligence, reachability and business context
Triage and SLAs	SLAs per Asset, risk acceptance with an owner and expiry date	Triage Engine rules, SLAs by severity or risk, approval workflow for accepted risk
Remediation	Jira with two-way sync, and GitHub Issues	Downstream Connectors for Azure DevOps, GitHub, GitLab, ServiceNow and more, plus the Remediation Center
Access and support	Local accounts and RBAC, community support on GitHub and the DefectDojo Community Slack	SSO, MFA, custom roles, instance-wide audit logs and commercial support with a dedicated customer success manager

SEE IT ON YOUR DATA

defectdojo.com/book-a-demo

A live session with your scanners, your findings and your questions. No canned deck.

COMPARE EDITIONS

defectdojo.com/pricing

See what each edition includes and get a quote sized to the findings volume you actually run.

07

Getting started with AI

The future of enterprise security is not hiring more people to manage more tools. It is letting AI handle the tactical noise so security leaders can spend their time on strategic defense.

01 **Consolidate and normalize**
Move from disconnected tools to one platform, and aggregate and deduplicate every finding in DefectDojo. Community Edition is a free place to start.

02 **Upgrade to Pro for enrichment and MCP**
DefectDojo Pro adds exploit intelligence, risk-based prioritization and a built-in MCP server, so your AI assistant works with clean, ranked data and your own permissions.

03 **Automate triage with a person in the loop**
Use Triage Engine rules for routine decisions and AI-assisted triage for the judgment calls, with every suggestion approved by a person.

04 **Add Sensei to ship the fix**
Turn ranked findings into pull requests your engineers can review and merge.

Conclusion

Whether the gain comes from better data quality or more efficient context, AI is a force multiplier that can give security operations a real performance boost.

The goal is for the hard work of normalizing and deduplicating data across hundreds of tools to become invisible. Then one engineer can manage security testing across global environments, and the team shifts from tactical data mining to strategic security architecture.

Join the underground.
Get your backlog to zero.

Security does not have to be a struggle, and your tools should not hold your data hostage. Bring your scanners and your backlog. defectdojo.com/book-a-demo

Produced by DefectDojo Inc. Questions: info@defectdojo.com or defectdojo.com. Sources: SANS 2025 survey; Help Net Security. IDC MarketScape: Worldwide Application Security Posture Management 2025. Product facts describe DefectDojo Pro as of October 2026; Sensei is available as an add-on to DefectDojo Pro and some capabilities require a Sensei license. All trademarks are the property of their respective owners. © 2026 DefectDojo Inc. All Rights Reserved.