DEFECTDOJO

BUYER'S GUIDE · 2026 EDITION
FOR SECURITY LEADERS

APPLICATION SECURITY POSTURE MANAGEMENT

THE ASPM BUYER'S GUIDE

A security leader's framework for evaluating Application Security Posture Management in 2026: the capabilities that matter, the architectures on offer, a weighted scorecard, the pitfalls, and a 30-day plan to decide with proof.

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

Contents

01 Executive summary PAGE 03

02 What is ASPM, and why now PAGE 04

03 The eight capabilities every ASPM must have PAGE 05

04 Architecture comparison PAGE 07

05 Evaluation scorecard PAGE 08

06 Common pitfalls PAGE 09

07 Why DefectDojo PAGE 10

08 Next steps: your 30-day evaluation plan PAGE 11

How to use this guide

Read sections 01 and 02 for the case you will make to leadership. Use section 03 to write your requirements, section 05 to score vendors on equal terms, and section 08 to run the evaluation itself in 30 days.

Every capability in section 03 ends with how DefectDojo Pro meets it. Those paragraphs are there so you can check the claims against a live instance, not so you can skip the scorecard.

WHO THIS IS FOR

CISOs, heads of application security and the engineers who will run the platform every day.

If you own a vulnerability backlog that spans more than one scanner, this guide applies to you whether you call the category ASPM, vulnerability management or something of your own.

KEY TAKEAWAY

The best ASPM platforms are tool-agnostic, collapse noise across every scanner, rank findings on real risk, close the loop with remediation, and never charge you more for securing more.

01

Executive summary

Application Security Posture Management has moved from an emerging category to an operational necessity. The question is no longer whether you need a layer above your scanners. It is which one.

Gartner projects that 80% of organizations in regulated industries that run application security testing will adopt ASPM by 2027, up from roughly 29% today. Frost & Sullivan expects the market to pass $1.7 billion by 2029, growing at about 30% a year.

The catalyst is plain. A modern enterprise runs dozens of scanners across code, dependencies, containers, cloud infrastructure and APIs. Each produces findings in its own format, on its own severity scale, with its own idea of what is critical. Without a unifying layer, teams drown in duplicates, cannot describe organizational risk in one sentence, and cannot enforce a remediation SLA they would defend to an auditor.

Two things changed in the last year. Attacks now move at AI speed, so the queue has to move at AI speed too. And AI-generated code has expanded the attack surface faster than any security team can hire. A platform that only stores findings is no longer enough. It has to rank them, route them and help close them.

This guide gives you a structured way to evaluate the options: the eight capabilities that matter, the architectural trade-offs, a weighted scorecard you can hand to procurement, the pitfalls that derail programs, and a four-week plan. Throughout, we show how DefectDojo Pro meets each requirement, drawing on the most widely deployed platform in security, with 50M+ downloads across 10,000+ organizations, and a Major Player rating from IDC in Application Security Posture Management.

FIVE THINGS TO INSIST ON

1. Tool-agnostic ingestion from every scanner you run, and the ones you will add.

2. Deduplication across every tool, not inside one, with root cause analysis so one fix closes many findings.

3. Prioritization on real risk. Severity is not risk.

4. Remediation that closes the loop: SLAs, tickets where engineers work, and an agent that ships the fix.

5. Pricing that never punishes coverage. Unlimited users and assets.

THE ONE-LINE TEST

Ask every vendor to load one week of real scan data from three of your tools, then show you the ten findings your team should fix first and why. If the answer takes more than a demo, the platform is not ready.

02

What is ASPM, and why now

ASPM analyzes security signals across development, deployment and operations to give you one normalized view of application risk, and to act on it. It is the orchestration layer that sits above individual scanners.

The problem it solves

Most enterprises operate between 10 and 50 distinct scanning tools: SAST, DAST, SCA, container and image scanners, cloud posture tools, secret scanners, pen test reports and bug bounty feeds. Each generates findings in a proprietary format with its own severity taxonomy. The result is a fragmented, noisy and often contradictory picture of risk.

- **Alert fatigue.** The same CVE appears in four scanners, and the team investigates it four times.
- **Inconsistent prioritization.** One tool says Critical, another says Medium, and the remediation effort follows whichever report was opened first.
- **No SLA enforcement.** Without one system of record you cannot measure time to remediation, let alone enforce it.
- **Compliance gaps.** Auditors want evidence. Scanners produce noise. Someone has to close the gap, and today it is usually a spreadsheet.

Market momentum

30%
ANNUAL GROWTH THROUGH 2029

$1.7B
PROJECTED MARKET SIZE, 2029

80%
REGULATED INDUSTRIES ADOPTING ASPM BY 2027

Sources: Gartner, Innovation Insight for Application Security Posture Management (2025); Frost & Sullivan, ASPM Market Report (2024 to 2029).

Three forces behind the acceleration

- **AI-generated code.** More code ships faster, written by tools that do not read your security policy. The attack surface grows every sprint.
- **Regulation with teeth.** The EU Cyber Resilience Act, NIS2, PCI DSS 4.0 and federal mandates all expect proof that vulnerabilities were found, tracked and resolved within policy.
- **Headcount does not scale with volume.** Finding volume grows with coverage. Security teams do not. Automation and orchestration are the only way to close the gap.

THE 2026 SHIFT

Attacks move at AI speed. An ASPM that stores and reports is a rear-view mirror. The platform you choose has to react at AI speed: rank on risk the moment a finding lands, route it, and put a fix in front of an engineer.

03

The eight capabilities every ASPM must have

Write your requirements from this list. Each capability states what to look for, then how DefectDojo Pro meets it, so you can verify the claim against a live instance during the proof of concept.

1 Universal tool integration

WHAT TO LOOK FOR

Your ASPM must ingest findings from every scanner in your environment without a custom parser for each one. Look for native support across SAST, DAST, SCA, container, cloud, secrets and pen test output, by file import and by API, plus scheduled pulls from the tools that expose one. Ask how long a new tool takes to onboard.

HOW DEFECTDOJO PRO DOES IT

500+ official integrations, one normalized data model. Findings arrive by file import, by API push from the pipeline, or through 130+ Connectors that pull on a schedule from tools such as Wiz, Snyk, Tenable, Qualys and CrowdStrike. The open parser architecture means a new tool is onboarded in days, not quarters.

2 Deduplication across every tool, with root cause

WHAT TO LOOK FOR

Running several scanners against the same code produces massive duplication. When 4,000 raw findings hold 40 unique issues and only 4 need action this week, the platform must collapse them automatically. Everyone dedupes inside one scanner. Insist on deduplication across every tool, configurable to the field, and on correlation that traces a cluster to its root cause.

HOW DEFECTDOJO PRO DOES IT

Deduplication runs across your entire stack and is configurable down to the field, so a CVE reported by four scanners is one finding with four sources. Root cause analysis traces the cluster to the change that closes it, so one fix closes many findings instead of one ticket per scanner.

3 Risk-based prioritization

WHAT TO LOOK FOR

Severity is not risk. Raw scanner scores ignore whether an exploit exists, whether the code path is reachable, and what the asset is worth. Your ASPM must enrich every finding with exploitability data such as EPSS and the CISA Known Exploited Vulnerabilities catalog, asset context and threat intelligence, then rank on the result. Ask to preview a scoring change before it lands.

HOW DEFECTDOJO PRO DOES IT

DefectDojo ranks on risk: EPSS, CISA KEV, live threat intelligence, asset context, reachability and many more signals, tunable per asset. Reachability analysis (beta) adds five verdicts, with KEV overriding the ceiling. The prioritization simulator previews scoring changes before they land, so the ranking is one you can defend to an auditor.

4 Triage automation

WHAT TO LOOK FOR

Most of the work in a backlog is not remediation. It is deciding what is a duplicate, what is fixed, what is accepted and who owns it. Look for a rules engine that closes fixed issues on re-scan, suppresses known false positives, applies risk acceptance policy and routes ownership without a human touching each finding.

HOW DEFECTDOJO PRO DOES IT

The Triage Engine runs auto-triage, auto-close and acceptance rules against every incoming finding. Fixed issues close on the next scan, accepted risk stays accepted with its evidence, and ownership follows your asset model. Human attention goes to the findings that need a decision.

5 SLA-driven remediation

WHAT TO LOOK FOR

Define remediation SLAs by severity, then measure them. The platform must track every finding against policy, warn before breach, escalate overdue findings and give leadership a compliance view without a spreadsheet.

HOW DEFECTDOJO PRO DOES IT

SLA enforcement with breach alerting: policies by severity, notifications before and at breach, and SLA compliance reporting for the executive view. The same clock feeds audit-ready reporting, so the evidence an assessor asks for already exists.

6 Workflow, ticketing and the pipeline

WHAT TO LOOK FOR

Findings must reach engineers where they already work. Evaluate bi-directional Jira and native ticketing for GitHub, GitLab, Azure DevOps and ServiceNow, not webhook workarounds. Then look at the pipeline: an API-first design, build and commit tracking on every scan, and the ability to gate a deployment on a threshold.

HOW DEFECTDOJO PRO DOES IT

Bi-directional Jira keeps ticket and finding in step. GitHub, GitLab, Azure DevOps and ServiceNow are native. Every scan records build ID, commit, branch, tag and repository, so a finding points at the code that introduced it, and the API-first architecture lets you gate releases on findings that matter.

7 Reporting, evidence and governance

WHAT TO LOOK FOR

Leadership needs risk trends, SLA compliance and remediation velocity. Practitioners need drill-down by asset, team and tool. Auditors need a system of record that proves what was found and when it was fixed. Look for an executive BI suite, a custom report builder, SSO with granular RBAC and a full audit trail.

HOW DEFECTDOJO PRO DOES IT

An executive BI suite and a custom report builder for the board, drill-down for the people doing the work, and audit-ready reports you hand to the assessor. SSO (SAML 2.0, OIDC), granular RBAC and a full audit trail are built in, alongside federal and compliance workflows: POA&M, CMMC posture, STIG checklists with CCI crosswalk, FIPS deployments and a PSIRT advisory workflow from intake to published advisory.

8 AI that closes findings

WHAT TO LOOK FOR

AI in an ASPM should do two things. It should let your own assistants query your posture through an open standard such as the Model Context Protocol, and it should take work off the queue. An AI that writes another ticket is not the same as one that ships the fix. Ask what the agent actually changes, and who controls it.

HOW DEFECTDOJO PRO DOES IT

Sensei is the specialized agent that ships the fix as a pull request for code findings and writes the misconfiguration correction for cloud posture findings. You set the autonomy: gate every change behind approval or let it run. DefectDojo Pro was also the first ASPM to support MCP, so any compatible model can work with your posture data.

WHY THE LAST ONE MATTERS MOST

Capabilities one through seven describe a system of record. The eighth turns it into a system of action. When the queue is bigger than the team, the only way to zero is a platform that reacts at AI speed and puts fixes in front of engineers instead of tickets.

04

Architecture comparison

Not every ASPM is built the same way. The architecture decides who controls your data, how far the platform can see, and what it costs to change your mind later.

CATEGORY	WHAT IT IS	STRENGTHS	LIMITATIONS
Open source foundations	Community-built platforms with a commercial edition on top. Transparent codebase, large contributor ecosystem.	No lock-in, code inspected by a global community, flexible deployment, visible roadmap.	The free edition needs internal resources to host and tune.
Scanner-native platforms	ASPM built by a scanner vendor as an add-on to its own portfolio.	Deep integration with the vendor's scanners, one bill.	Bias toward the vendor's tools, thin third-party support, lock-in risk when you change scanners.
Cloud-only ASPM	SaaS-only platforms built for cloud-first organizations.	Fast deployment, automatic updates, minimal infrastructure.	Data residency constraints, no self-hosted option, less control over the environment.
GRC add-ons	ASPM features bolted onto a governance, risk and compliance suite.	One platform for GRC and security, an existing enterprise contract.	ASPM is not the core focus; deduplication, triage and remediation depth suffer.

WHERE DEFECTDOJO FITS

DefectDojo spans the first and third categories. It began as an OWASP flagship open source project, now at 50M+ downloads, so the code is transparent and there is no lock-in. DefectDojo Pro adds the risk engine, cross-stack deduplication, the Triage Engine, Sensei remediation, SSO and RBAC, and runs as SaaS or self-hosted at the same price. Your data belongs to you in either case.

Questions to ask about architecture

- Can I run it in my own environment, and does that cost more?
- If I replace a scanner next year, what breaks?
- Can I export every finding, with its history, without a services engagement?
- Is the core inspected by anyone outside the vendor?
- Does the platform have a written data residency answer for every region I operate in?

05

Evaluation scorecard

Score each vendor from 1 to 5 on every criterion, multiply by the weight, and compare totals. Adjust the weights to your program before the first demo, not after.

CAPABILITY	WEIGHT	WHAT TO EVALUATE	VENDOR A	VENDOR B	VENDOR C
Tool integration breadth	20%	Native parsers and Connectors, effort to add a custom source, time to onboard a new tool			
Deduplication accuracy	15%	Cross-tool dedupe, field-level configuration, root cause correlation, false-positive rate			
Risk-based prioritization	15%	EPSS and KEV enrichment, threat intelligence, asset context, reachability, scoring preview			
Triage automation and SLAs	10%	Rules engine, auto-close on re-scan, acceptance policy, SLA enforcement and breach alerting			
Workflow and ticketing	10%	Bi-directional Jira, native GitHub, GitLab, Azure DevOps and ServiceNow			
Reporting and compliance	10%	Executive BI, custom reports, audit-ready evidence, SSO, RBAC, audit trail			
Pipeline and API	5%	API-first design, build and commit tracking, deployment gating			
AI remediation and MCP	5%	Does the agent ship fixes, who controls autonomy, MCP support			
Scalability	5%	Concurrent scan handling, performance at millions of findings			
Pricing and licensing	5%	Transparent model, unlimited users and assets, no per-scanner or per-asset penalty			
Total weighted score	100%				

RUN IT ON YOUR DATA

Run the proof of concept with your top two or three vendors using real scan data from your environment. Synthetic demos hide the integration pain that only surfaces at production finding volumes. Score the POC, not the demo.

WEIGHT THE LAST ROW HONESTLY

Pricing is 5% of the score and 100% of the renewal conversation. Ask for the pricing model in writing before the POC starts, and check whether adding users, assets or scanners changes the number.

06

Common pitfalls

Most failed ASPM programs did not pick a bad product. They picked the wrong shape of product for their scanners, their engineers or their budget. Six ways that happens, and how to avoid each one.

1 Choosing a scanner vendor's ASPM

When the same vendor sells the scanners and the layer above them, third-party tools become second-class citizens. Integration depth varies, and swapping a scanner later means fighting your own platform. Choose a platform that treats every scanner equally.

2 Ignoring deduplication quality

Some platforms advertise deduplication and ship exact-match logic that misses the same issue reported two ways. Ask for a live demonstration with your scan data from several tools on one application. The difference between 60% and 95% accuracy is measured in engineer hours every week.

3 Per-user, per-asset or per-scanner pricing

These models penalize the behavior you want: more coverage, more people looking, more tools. Costs balloon as the program matures. Insist on unlimited users and assets, and on a price tied to the work the platform does for you.

4 Underestimating migration

Moving years of findings into a new platform is real work. Evaluate import coverage, API documentation and the quality of professional services. A platform with 500+ native integrations removes most of the parser writing from the plan.

5 Buying features you cannot operationalize

Advanced analytics and AI are powerful only if the team can run them. Start with the fundamentals: one queue, deduplication, prioritization, SLAs and ticketing. Layer on the rest as the program matures, and pick a platform where the layers are already there.

6 Neglecting the developer experience

If engineers hate the tool, adoption stalls. Check how findings show up in their workflow: a Jira ticket with a clear fix, a pull request they can review, a false positive they can dismiss without leaving their tools. The best ASPM removes friction for developers, not only for security.

A SEVENTH, FOR 2026

Mistaking an AI that describes findings for an AI that fixes them. Summaries and chat over a backlog are useful. They are not remediation. Ask every vendor what their agent changes in your repositories and cloud accounts, and who approves it.

07

Why the world runs on DefectDojo

Open source since 2014, born inside a real application security program, built by 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, into one normalized model. **Everyone dedupes inside one scanner. Dojo dedupes across every tool.**

Sensei ships the fix

The specialized agent to solve security. Code fixes arrive as pull requests, cloud misconfigurations as corrections. **You set the autonomy. Control stays yours.**

Pricing that never punishes coverage

Unlimited users and assets. You pay for the findings the platform processes, pay as you go or pre-pay and save 48% or more. **SaaS or self-hosted, same price, and your data belongs to you.**

Risk ranking you can defend to an auditor

EPSS, CISA KEV, threat intelligence, asset context, reachability and many more signals, tunable per asset, with a simulator to preview changes. **Severity is not risk.**

Evidence, not spreadsheets

Audit-ready reporting, SLA enforcement, SSO, RBAC and a full audit trail. Federal and compliance workflows built in: **POA&M, CMMC posture, STIG with CCI crosswalk, FIPS deployments, PSIRT advisories.**

Prove it before you buy it

Community Edition is the full open source platform, in production at 10,000+ organizations. **Run the free one in your stack, then bring the same system of record to Pro.**

CAPABILITY	COMMUNITY EDITION · WHERE PROGRAMS START	DEFECTDOJO PRO · WHERE PROGRAMS SCALE
Deduplication	Within a scanner, limited across	Across your entire stack, configurable to the field, with root cause analysis
Prioritization	Severity and CVSS from the scanner	Risk engine with EPSS, CISA KEV and threat intel, tunable per asset, with a simulator
Ingestion	All 500+ integrations by file import or API push	Adds 130+ Connectors that pull on a schedule
Triage and SLAs	Manual triage, basic SLA tracking	Triage Engine rules, SLA enforcement with breach alerting
Remediation	Bi-directional Jira	Adds GitHub, GitLab, Azure DevOps, ServiceNow, and Sensei shipping the fix
Governance	Community support on OWASP Slack and GitHub	SSO, granular RBAC, full audit trail, executive BI, SLA-backed support

SEE IT ON YOUR DATA

defectdojo.com/book-a-demo

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

PRICING IN THE OPEN

defectdojo.com/pricing

Every tier, every inclusion and the pay as you go rate, published. Size it yourself before you talk to anyone.

08

Next steps: your 30-day evaluation plan

A structured evaluation prevents analysis paralysis and compares vendors on equal footing. Adapt the four weeks to your procurement calendar; keep the order.

WEEK	ACTIVITY	DELIVERABLE
Week 1	Define requirements and success criteria. Inventory your scanners, ticketing systems and pipelines. Pick two or three representative applications and a set of real scan data from at least three tools for the POC.	Requirements document and the scorecard from section 05 with your weights filled in
Week 2	Run vendor demos with your requirements as the agenda. Insist on your scan data, not a canned environment. Evaluate deduplication quality, setup time per integration and how the ranked queue reads.	Completed scorecards for every vendor and a shortlist of two or three
Week 3	Run the proof of concept with the shortlist. Import the real data, configure SLAs and ticketing, measure dedupe accuracy, and let an engineer receive and close a finding end to end. Ask the AI to fix something.	POC results: dedupe accuracy, integration time, time from finding to merged fix
Week 4	Evaluate pricing models and confirm the deployment model, data residency and support commitments. Check what happens to the price when users, assets and scanners grow. Build the business case.	Vendor recommendation with business case and ROI projection for leadership sign-off

What a good POC proves

- Three of your tools, one queue, no duplicates.
- A ranked list your team agrees with, and can explain.
- An SLA breach that alerts the right person.
- A finding that became a ticket, then a pull request, then closed on the next scan.

START WITH THE FREE ONE

Community Edition is the full open source platform. Prove the integrations and the deduplication in your own stack this week, then bring the same system of record to DefectDojo Pro for risk ranking, the Triage Engine, Sensei remediation, SSO and enterprise support.

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. Market data from Gartner, Innovation Insight for Application Security Posture Management (2025), and Frost & Sullivan, ASPM Market Report (2024 to 2029). IDC MarketScape vendor assessment for Application Security Posture Management. Product facts describe DefectDojo Pro as of September 2026; reachability analysis is in beta. All trademarks are the property of their respective owners. © 2026 DefectDojo Inc. All Rights Reserved.