DEFECTDOJO

BUYER'S GUIDE · 2026 EDITION
FOR SECURITY LEADERS

RISK-BASED VULNERABILITY MANAGEMENT

THE RBVM BUYER'S GUIDE

A security leader's framework for moving beyond CVSS in 2026: the risk inputs that matter, the capabilities to require, a weighted scorecard, the pitfalls, and a 30-day plan to find the 3% of vulnerabilities that pose real risk.

500+
OFFICIAL INTEGRATIONS

50M+
DOWNLOADS

10,000+
ORGANIZATIONS

IDC
MAJOR PLAYER IN ASPM

Contents

01 Executive summary PAGE 03

02 Why CVSS alone fails PAGE 04

03 The six risk inputs that drive real prioritization PAGE 05

04 The eight capabilities every RBVM platform must have PAGE 06

05 Architecture comparison PAGE 08

06 Evaluation scorecard PAGE 09

07 Common pitfalls PAGE 10

08 Why DefectDojo PAGE 11

09 Next steps: your 30-day evaluation plan PAGE 12

How to use this guide

Read sections 01 and 02 for the case you will make to leadership. Use sections 03 and 04 to write your requirements, section 06 to score vendors on equal terms, and section 09 to run the evaluation itself in 30 days.

Every risk input and capability ends with how DefectDojo Pro handles it. Those notes are there so you can check the claims against a live instance, not so you can skip the scorecard.

WHO THIS IS FOR

CISOs, vulnerability management leads and the engineers who decide what gets fixed first.

If your backlog is ordered by CVSS and still growing, this guide applies to you, whatever your scanners and whatever you call the program.

KEY TAKEAWAY

Only about 3% of vulnerabilities pose real, exploitable risk. The right RBVM platform finds them automatically, shows why they rank where they do, and moves them to the people who fix them.

01

Executive summary

The challenge is no longer finding vulnerabilities. It is knowing which ones to fix first, and being able to prove the order was right.

CVE issuance rose about 20% year over year in 2025, and the average time to exploit a new vulnerability has dropped to five days. Yet in any given environment only about 3% of vulnerabilities represent real, exploitable risk. A queue sorted by CVSS hides that 3% under thousands of findings marked Critical.

CVSS measures theoretical severity. It says nothing about whether an exploit exists, whether the vulnerable code is reachable, or what the affected asset is worth.

Risk-Based Vulnerability Management (RBVM) combines vulnerability data with threat intelligence such as CISA KEV and EPSS, asset context such as business criticality and internet exposure, and environmental factors such as reachability. The output is one backlog, ranked by actual risk. Gartner places this prioritization at the center of its Continuous Threat Exposure Management (CTEM) framework.

This guide gives you a structured way to evaluate RBVM platforms: the six risk inputs that matter, the eight capabilities to require, the architectures on offer, a weighted scorecard and a four-week plan. Throughout, we show how DefectDojo Pro meets each requirement.

FIVE THINGS TO INSIST ON

1. KEV and EPSS on every finding, refreshed automatically.

2. A risk score you can read, weight and defend.

3. Business context at the asset level, feeding the rank.

4. Deduplication across every tool before anything is ranked.

5. A ranked list that turns into tickets, SLAs and fixes.

THE ONE-LINE TEST

Load the same week of scan data into every platform you evaluate, then ask your two most senior engineers to pick their top ten by hand. The platform whose top ten matches theirs, and can explain why, is doing real prioritization.

THE 2026 SHIFT

Attacks now move at AI speed, and AI-generated code ships faster than any team can review it. A platform that only stores and reports findings is a rear-view mirror. It has to rank each finding the moment it lands, route it to an owner, and put a fix in front of an engineer.

02

Why CVSS alone fails

CVSS has rated vulnerability severity since 2005. It is a useful baseline. It was never designed to decide what your team fixes on Monday.

3%
OF VULNERABILITIES POSE REAL, EXPLOITABLE RISK

5 days
AVERAGE TIME TO EXPLOIT A NEW VULNERABILITY

18%
OF ORGANIZATIONS USE EPSS TO PRIORITIZE

- **It measures theoretical severity, not your risk.** A CVSS 9.8 in an air-gapped internal system with no sensitive data poses far less risk than a CVSS 7.2 in an internet-facing payment service. CVSS cannot see the difference because it cannot see your environment.
- **It does not move when attackers do.** CVSS scores are static. A 6.5 with a public exploit in use by ransomware groups is more urgent than a 9.0 with no known exploit. KEV and EPSS capture that shift. CVSS does not.
- **It floods the queue with Critical.** When every scanner flags thousands of Critical and High findings, everything is a priority, which means nothing is. Teams fall back on first come, first served, the least effective order there is.
- **It ignores business context.** The same vulnerability means different things in a revenue-generating production app and in a development sandbox. Criticality, data sensitivity, regulatory scope and revenue all belong in the rank.

CVSS only vs. risk-based

FACTOR	CVSS ONLY	RISK-BASED (RBVM)
Severity source	Static CVSS score	CVSS plus KEV, EPSS and business context
Exploit awareness	None	Active exploitation (KEV) and predicted exploitation (EPSS)
Asset context	None	Criticality, internet exposure, data sensitivity
Output	Static severity list	A ranked backlog that moves as the threat moves
Remediation effort	Spent chasing volume	Spent on the 3% that matters

Sources: FIRST.org EPSS v4; CISA Known Exploited Vulnerabilities catalog; Seemplicity 2025 Remediation Operations Report.

03

The six risk inputs that drive real prioritization

Risk-based prioritization is only as good as the signals feeding it. A mature RBVM platform combines all six into one score for each finding.

1 Known Exploited Vulnerabilities (KEV)

WHAT IT TELLS YOU

CISA's KEV catalog lists CVEs confirmed as actively exploited in the wild. A finding on the KEV list is urgent whatever its CVSS score says.

HOW DEFECTDOJO PRO USES IT

Every finding with a CVE is checked against the KEV catalog, including known ransomware use, and the result feeds its priority score. SLA deadlines never run past the KEV due date.

2 Exploit Prediction Scoring System (EPSS)

WHAT IT TELLS YOU

EPSS estimates the probability, from 0 to 1, that a vulnerability will be exploited in the next 30 days. EPSS v4, released in March 2025, improved its accuracy, yet only 18% of organizations use it.

HOW DEFECTDOJO PRO USES IT

EPSS scores are looked up for every finding with a CVE, whichever tool reported it, and refreshed every 24 hours. Air-gapped installs load them from a file.

3 Asset criticality and business context

WHAT IT TELLS YOU

A vulnerability in a customer-facing production system carries different risk than the same one in an internal sandbox. Criticality, data classification, regulatory scope and revenue belong at the asset level.

HOW DEFECTDOJO PRO USES IT

Business attributes on each Asset, such as criticality, internet exposure, revenue and user records, feed the priority score directly. SLAs can be assigned automatically by business criticality.

4 Exploitability and reachability

WHAT IT TELLS YOU

Can an attacker reach the vulnerable code? Is there a public exploit? Is the component internet-facing or isolated behind several network layers? The answers change real risk dramatically.

HOW DEFECTDOJO PRO USES IT

Pro shows whether vulnerable code is actually called and whether a host can be reached from the internet, and each verdict shows where it came from. Exploit-maturity data is added to every finding with a CVE.

5 Compliance and regulatory exposure

WHAT IT TELLS YOU

A vulnerability in a system that handles cardholder data, health records or EU personal data carries regulatory risk on top of technical risk: fines, failed audits and mandatory disclosure.

HOW DEFECTDOJO PRO USES IT

Findings map to the controls of nine bundled frameworks, including PCI DSS 4.0.1, SOC 2, NIS2 and the EU Cyber Resilience Act, and SLAs can follow regulatory deadlines such as FedRAMP response times.

6 Threat changes over time

WHAT IT TELLS YOU

Has exploit activity increased? Has a new proof of concept been published? Priorities should update as the threat changes, not only when the next scan runs.

HOW DEFECTDOJO PRO USES IT

Exploit intelligence refreshes every 24 hours, so a finding added to KEV overnight rises in the queue the next day, and a Triage Engine rule can act on the change as it happens.

04

The eight capabilities every RBVM platform 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 Automated KEV and EPSS enrichment

WHAT TO LOOK FOR

Every finding with a CVE must arrive enriched with current KEV status and EPSS probability. No manual lookups, no CSV imports, and the data must refresh on its own.

HOW DEFECTDOJO PRO DOES IT

EPSS, KEV and exploit-maturity data are looked up for every finding with a CVE and updated every 24 hours, whichever scanner found it.

2 A risk score you can configure

WHAT TO LOOK FOR

A black-box score is not enough. You must see which inputs drive the score, change their weights to match your risk appetite, and weight different kinds of assets differently.

HOW DEFECTDOJO PRO DOES IT

Every finding gets a priority score and a risk band: Urgent, Needs Action, Medium or Low. Weights are customizable, and each Asset can run its own Prioritization Engine.

3 Business context integration

WHAT TO LOOK FOR

Define criticality, data classification, regulatory scope and revenue per asset, and have the platform factor them into the rank automatically.

HOW DEFECTDOJO PRO DOES IT

Business attributes recorded on each Asset feed the priority score. Asset Hierarchy lets parent and child Assets reflect how the business is really built.

4 Cross-tool deduplication

WHAT TO LOOK FOR

When several scanners report the same vulnerability, the platform must collapse them into one finding before ranking, not create separate entries that inflate the backlog.

HOW DEFECTDOJO PRO DOES IT

Cross-tool deduplication is chosen per tool in the UI, and Root Causes groups the findings one fix would clear, ranked by how many findings and Assets each cause affects.

5 SLA-driven remediation

WHAT TO LOOK FOR

Set remediation deadlines by risk, not only by CVSS severity. Notify owners before a breach, escalate overdue findings, and report SLA compliance without a spreadsheet.

HOW DEFECTDOJO PRO DOES IT

SLAs by severity or by risk, with a choice of when the clock starts: at detection, at publication of the vulnerability, or when a fix is available. Deadlines never exceed the CISA KEV due date.

6 Triage automation

WHAT TO LOOK FOR

Routine decisions should not need a human: escalate KEV findings, assign owners, tag compliance scope, and close what a rescan proves fixed.

HOW DEFECTDOJO PRO DOES IT

The Triage Engine is a no-code rules engine. Rules run when a scan lands or a finding changes, then update findings, assign owners, accept risk, open tickets or send messages. New rules start in Simulate mode.

7 Priorities that update with the threat

WHAT TO LOOK FOR

The rank must change when the threat does: a new KEV listing, an EPSS shift, a new public exploit. Waiting for the next scheduled scan is too slow.

HOW DEFECTDOJO PRO DOES IT

Exploit intelligence refreshes every 24 hours across every finding with a CVE, and Triage Engine rules can trigger on the change to escalate, notify or open a ticket.

8 Remediation that closes the loop

WHAT TO LOOK FOR

A ranked list is useless if engineers cannot act on it. Look for concrete fix guidance, tickets in the tools engineers already use, and help shipping the fix itself.

HOW DEFECTDOJO PRO DOES IT

The Remediation Center turns open findings into a short list of fixes, each due by the earliest SLA it clears. Downstream Connectors push work to Jira, Azure DevOps, GitHub, GitLab, ServiceNow and more. Sensei, an add-on, opens the pull request.

WHY THE LAST ONE MATTERS MOST

Capabilities one through seven decide what to fix. The eighth decides whether it gets fixed. When the queue is bigger than the team, the only way toward zero is a platform that puts fixes in front of engineers, not more tickets.

05

Architecture comparison

RBVM ships in very different architectures. The architecture decides how far the platform can see, who controls your data, and what it costs to change your mind later.

CATEGORY	WHAT IT IS	STRENGTHS	LIMITATIONS
Open source foundations	Community-built vulnerability management platforms with a commercial edition that adds RBVM.	Transparent code, tool-agnostic, no lock-in, large integration ecosystem.	The free edition orders findings by severity; full risk-based ranking comes with the commercial tier.
Scanner-native RBVM	Prioritization built into a scanning vendor's platform, using its own scanner data as the main input.	Deep integration with the vendor's scanners, one platform for scanning and ranking.	Bias toward the vendor's own data, thin third-party support, lock-in risk.
Standalone RBVM	Purpose-built tools that aggregate several scanners for risk-based ranking.	Deep prioritization logic, strong threat intelligence.	Another tool to integrate; often no ticketing, SLAs or compliance evidence.
Exposure management	Broader platforms that extend RBVM into attack surface management and risk quantification.	Wider risk view, attack path analysis, asset discovery.	Higher complexity and cost; may exceed current program maturity.

WHERE DEFECTDOJO FITS

DefectDojo starts from an open source foundation: an OWASP flagship project with 50M+ downloads and 500+ integrations. DefectDojo Pro adds full RBVM, with KEV, EPSS, exploitability, reachability and business context in one score, plus the lifecycle that standalone RBVM tools lack: SLAs, ticketing, the Triage Engine and compliance evidence.

Questions to ask about architecture

- Can I see, and change, every input to the risk score?
- If I replace a scanner next year, does my ranking still work?
- Does the ranked list become tickets and SLAs without another tool?
- Can I run it in my own environment, including air-gapped?
- Can I export every finding with its history, without a services engagement?

06

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
KEV and EPSS enrichment	15%	Automated sync, refresh frequency, automatic escalation of KEV findings			
Risk scoring configurability	15%	Adjustable weights, custom inputs, transparent scoring logic			
Business context integration	15%	Asset-level criticality, data classification, regulatory tagging			
Cross-tool deduplication	10%	Accuracy across scanners, tunable matching, root cause grouping			
SLA-driven remediation	10%	Risk-based SLA policies, notification and escalation, compliance reporting			
Triage automation	10%	Escalate, assign, tag and close by rule, with a safe way to test rules			
Scanner integration breadth	10%	Native parsers, scheduled connectors, API ingestion, SARIF and CycloneDX			
Dynamic priority updates	5%	Continuous EPSS and KEV refresh, recalculated priorities			
Remediation guidance	5%	Fix lists, advisory links and context in tickets, AI-assisted fixes			
Pricing and licensing	5%	Transparent model, unlimited users and assets, no per-scanner penalty			
Total weighted score	100%				

RUN IT ON YOUR DATA

Import the same scan data into every vendor during the POC and compare the ranked output. Ask whether the top ten matches what your senior engineers would pick by hand. That is the real test of a risk scoring engine.

WEIGHT THE LAST ROW HONESTLY

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

07

Common pitfalls

Most RBVM programs that stall did not buy a bad product. They bought a ranking without a workflow, or a workflow without a ranking they could trust. Six ways that happens.

1 Treating CVSS as good enough

CVSS measures theoretical severity. Teams that prioritize on CVSS alone spend more than 60% of remediation effort on findings that pose no practical threat. RBVM is the difference between chasing volume and reducing risk.

2 Adopting EPSS without KEV

EPSS predicts exploitation; KEV confirms it. When a vulnerability is on the KEV list, that fact should override its EPSS score. Use both: KEV as the fix-now signal, EPSS as the early warning.

3 Black-box risk scoring

If you cannot see or configure how a vendor calculates priority, you cannot trust it or defend it to an auditor. Ask what the inputs are, how they are weighted, and whether you can change the weights.

4 RBVM without lifecycle management

A ranked list does nothing on its own. If the tool cannot push findings to Jira, enforce SLAs and track remediation, you will need a second platform to act on its output.

5 Ignoring business context

A KEV-listed vulnerability on a decommissioned test server is less urgent than a non-KEV finding in production payments. Make sure asset-level business context is part of the score.

6 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.

A SEVENTH, FOR 2026

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

08

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.

09

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	Decide which risk inputs matter most to your organization. Inventory your scanners, pick two or three test applications, and prepare scan data from several tools for the POC.	Requirements document and the scorecard from section 06 with your weights
Week 2	Run vendor demos on your data. Compare how each platform ranks the same findings, and whether its top ten matches what your senior engineers would choose.	Completed scorecards and a shortlist of two or three
Week 3	Run the POC. Import multi-scanner data, set business context per asset, check KEV and EPSS coverage, configure SLA policies, and push one finding through ticketing to a fix.	POC results: ranking accuracy, enrichment coverage, SLA setup effort
Week 4	Evaluate pricing and confirm the deployment model, data residency and support commitments. Build the business case around the remediation effort you will save.	Vendor recommendation with business case and projected ROI

What a good POC proves

- KEV and EPSS on every finding with a CVE, without a single CSV.
- A top ten your senior engineers agree with, and can explain.
- An SLA set by risk that alerts the right person before it breaches.
- A ranked finding that became a ticket, then a fix, then closed on the next scan.

START WITH THE FREE ONE

Community Edition is the free open source platform. Prove the integrations in your own stack this week, then move the same data to DefectDojo Pro for risk-based prioritization, the Triage Engine, SLAs by risk 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 the Gartner CTEM framework, FIRST.org EPSS v4, the CISA KEV catalog, the Seemplicity 2025 Remediation Operations Report and publicly available vendor documentation. IDC MarketScape vendor assessment for Application Security Posture Management. Product facts describe DefectDojo Pro as of October 2026; Sensei is available as an add-on to DefectDojo Pro. All trademarks are the property of their respective owners. © 2026 DefectDojo Inc. All Rights Reserved.