THE CYBER TOOLCHAIN

NO. 009

Head

THE WEEK'S TOP PICKS

PUBLISHED AUGUST 28, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

The best tooling updates that shipped this week.

// HOW THESE PICKS ARE MADE

Every feature release from the 902 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and choose our top picks, judged on two questions: **how deep and complete is the single best capability in the release**, and **how much it changes what you can actually do**.

A tool is judged on everything it shipped that week, so a project that releases daily gets credit for the sum — and still only takes one slot. We would rather run a short list than a padded one.

// THE PICKS

01	**SonarQube**	SAST / SCA / Code Security
02	**Aikido Security**	SAST / SCA / Code Security
03	**Zscaler**	Secure Access (VPN / SASE / ZTNA)
04	**JFrog**	SAST / SCA / Code Security
05	**defectdojo**	Vulnerability Management

01 SonarQube

1 RELEASE · 2026-08-27

SAST / SCA / Code Security

Trusted by 7M+ developers, SonarQube is the AI code review and code verification platform that catches quality and security issues before they reach production.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

SonarQube's Hunter Agent is now GA on SonarQube Cloud, autonomously finding broken access control (IDOR, missing authorization, privilege escalation), business logic flaws, and auth/session weaknesses that SAST structurally cannot evaluate — validating each finding for exploitability before surfacing it as a normal SonarQube issue at 80–90% precision.

> SonarQube launched Hunter Agent, a GA AI security agent on SonarQube Cloud's Enterprise plan that finds broken access control, business logic, and authentication flaws that traditional SAST cannot evaluate.

└──▷ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

01 Hunter Agent for logic and access-control flaws NEW

81

Hunter Agent is a new AI security agent, GA on SonarQube Cloud (Enterprise plan), that detects broken access control (IDOR, missing authorization checks, privilege escalation, CSRF), business logic flaws (skipped workflow steps, missing rate limits), and authentication/session management flaws (session fixation, non-expiring sessions, weak password recovery, missing MFA) — vulnerability classes SAST cannot structurally evaluate. Every candidate finding is validated for exploitability by a dedicated subagent before surfacing, achieving 80–90% average precision, and each finding carries an auditable discovery path showing how the agent reached its conclusion. Confirmed findings appear directly as SonarQube issues in the existing triage workflow with no separate portal or install, can be triggered on-demand or scheduled weekly/monthly against a main branch, and are delivered asynchronously. It scales to codebases beyond 10 million lines via a structured multi-phase pipeline (Analyze, Explore, Validate, Synthesize) with incremental scanning that cuts large-repo re-scan time roughly in half, and works across the 40+ languages already supported by SonarQube with no additional configuration.

— Rich mechanism and numbers given, but no exact command or config key.

launch-20260827-9c110bbd

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Aikido Security

10 RELEASES · 2026-08-25 → 2026-08-26

SAST / SCA / Code Security

Aikido Security is a developer-focused security platform that scans code, dependencies, containers, and cloud infrastructure to detect and prioritize vulnerabilities for fast remediation.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

Agentic Android (APK) pentesting installs a real build on a device and attacks it over ADB, with every finding exploit-confirmed and reproduction steps given as exact ADB commands — and it assesses the app plus its backend API in one run. The same week also shipped agentic Code Audit for logic flaws, code-based DSPM data-flow tracing, and real-time malware blocking on developer devices.

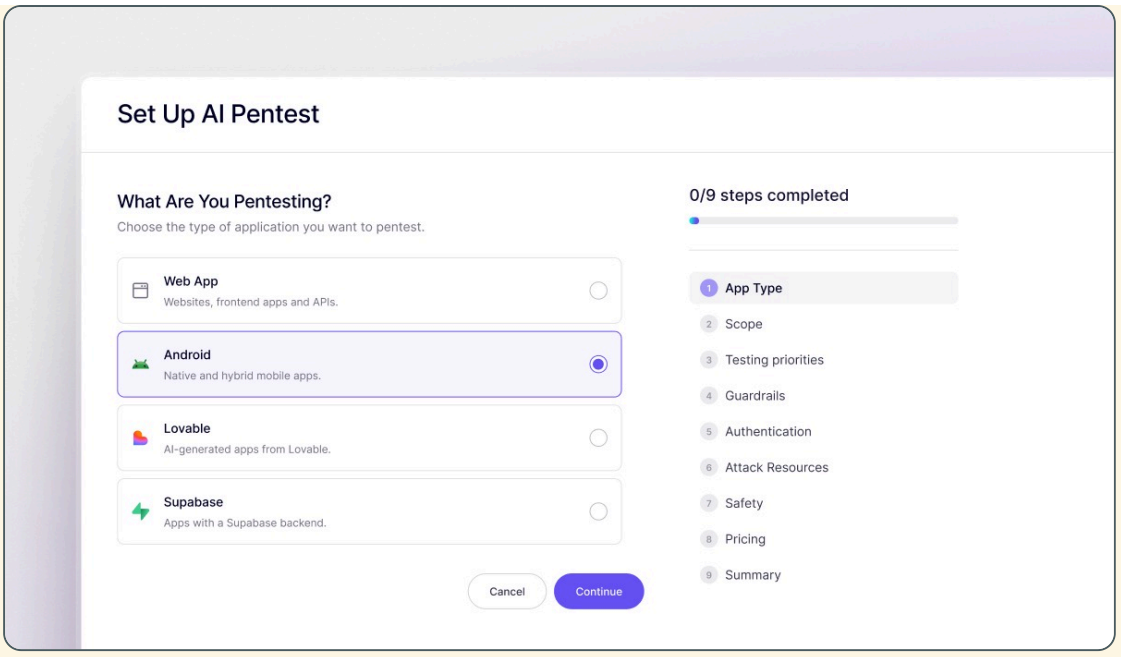
> Aikido's biggest addition this window is agentic Android (APK) pentesting with ADB-based exploit confirmation, alongside a wave of new AI agents: a Device Protection agent for company devices (now including Linux), LLM-powered Secrets AutoTriage and CVE Exploitability Analysis, code-based Data Exposure Audit (DSPM), and pentest-grade Code Audit that needs no live environment. It also took GCP VM scanning to general availability and shipped smaller platform, reporting and API updates.

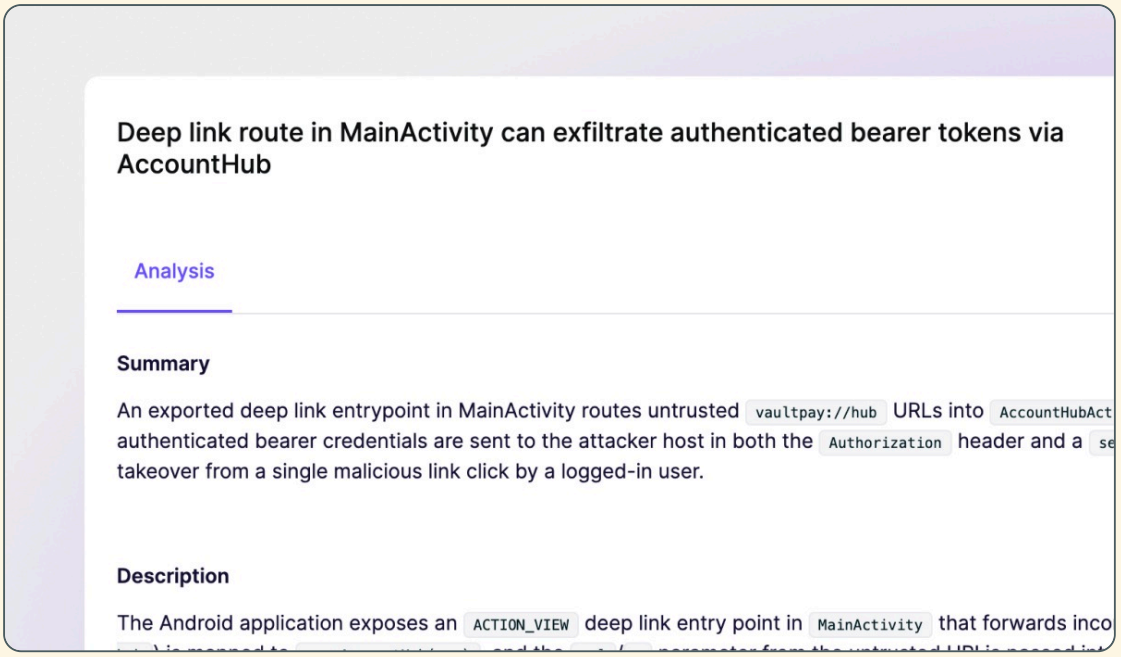
└──▷ WHAT SHIPPED · 13 FEATURES most completely described first ? what's the number?

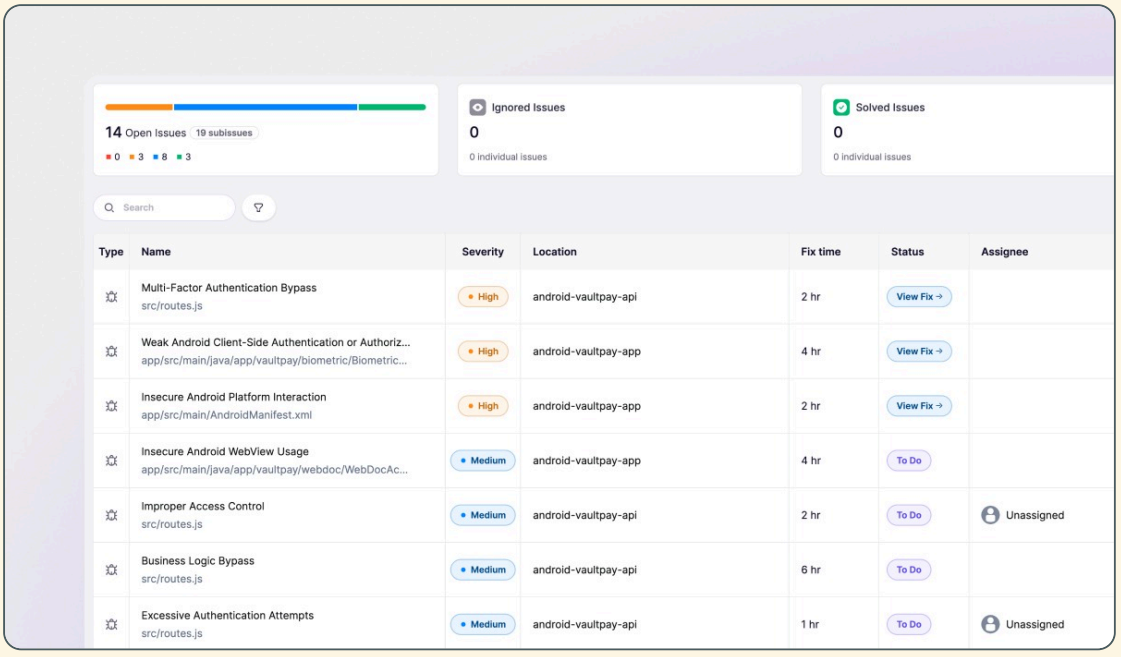
01 Agentic Android (APK) pentesting NEW

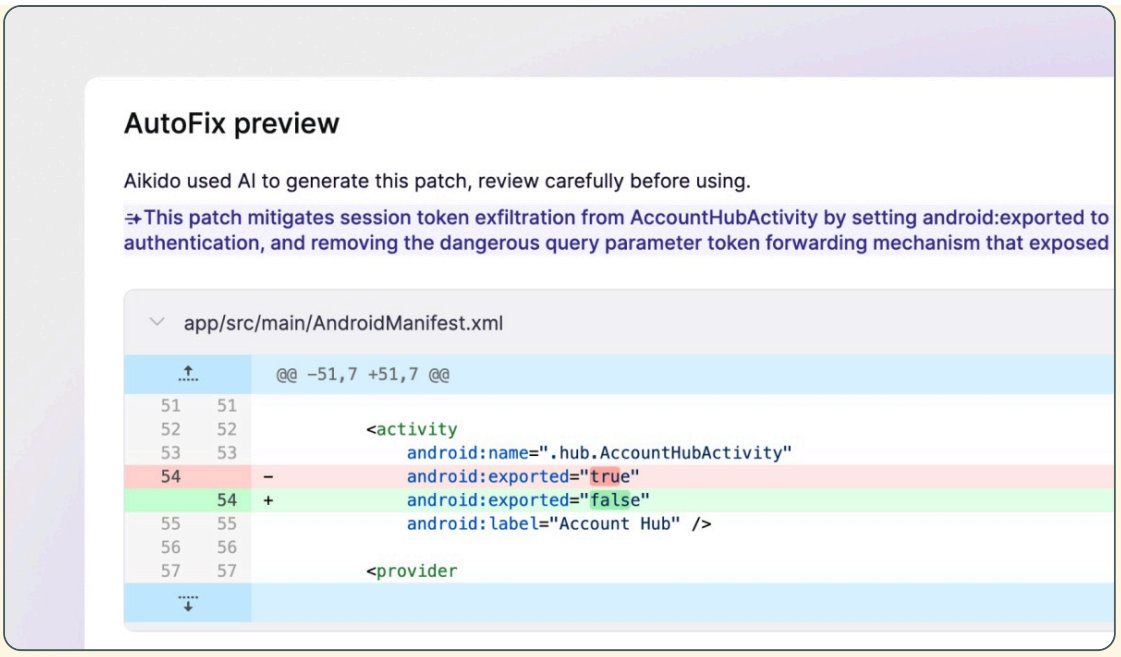
88

Aikido agents install a real APK on a test device, log in as a user, and attack the app over ADB, producing exploit-confirmed findings with the exact ADB commands needed to reproduce each issue. Runs can simultaneously assess the Android app and its backend API in one report, cover unsafe WebView usage, deep-link and exported-component abuse, client-side authentication issues, injection flaws, broken access control, weak session handling, unsafe API behavior, IDORs, cross-tenant access, and authentication bypasses, and produce a recon report and attack plan with diagrams, threat scenarios, and the full agent reasoning trace (including attempts that weren't confirmed). Assessments are 'rightsized' to app size rather than flat-fee, are currently white-box only (source repo access required alongside the APK, black-box support coming soon), and require certificate pinning, root detection, emulator detection, and RASP to be disabled in the uploaded build.


Set Up AI Pentest
What Are You Pentesting?
Choose the type of application you want to pentest.
Web App
Websites, frontend apps and APIs.
Android
Native and hybrid mobile apps.
Lovable
AI-generated apps from Lovable.
Supabase
Apps with a Supabase backend.
Cancel
Continue
0/9 steps completed
1 App Type
2 Scope
3 Testing priorities
4 Guardrails
5 Authentication
6 Attack Resources
7 Safety
8 Pricing
9 Summary


Deep link route in MainActivity can exfiltrate authenticated bearer tokens via AccountHub
Analysis
Summary
An exported deep link entrypoint in MainActivity routes untrusted vaultpay://hub URLs into AccountHubAct
authenticated bearer credentials are sent to the attacker host in both the Authorization header and a se
takeover from a single malicious link click by a logged-in user.
Description
The Android application exposes an ACTION_VIEW deep link entry point in MainActivity that forwards inco


14 Open Issues 19 subissues
0 3 8 3
Ignored Issues
0
0 individual issues
Solved Issues
0
0 individual issues
Search
Type Name Severity Location Fix time Status Assignee
Multi-Factor Authentication Bypass
src/routes.js
High
android-vaultpay-api
2 hr
View Fix
Weak Android Client-Side Authentication or Authoriz...
app/src/main/java/app/vaultpay/biometric/Biometric...
High
android-vaultpay-app
4 hr
View Fix
Insecure Android Platform Interaction
app/src/main/AndroidManifest.xml
High
android-vaultpay-app
2 hr
View Fix
Insecure Android WebView Usage
app/src/main/java/app/vaultpay/webdoc/WebDocAc...
Medium
android-vaultpay-app
4 hr
To Do
Improper Access Control
src/routes.js
Medium
android-vaultpay-api
2 hr
To Do
Unassigned
Business Logic Bypass
src/routes.js
Medium
android-vaultpay-api
6 hr
To Do
Excessive Authentication Attempts
src/routes.js
Medium
android-vaultpay-api
1 hr
To Do
Unassigned



After merging an AutoFix PR for a confirmed Android finding, retest against the patched build without rescoping.

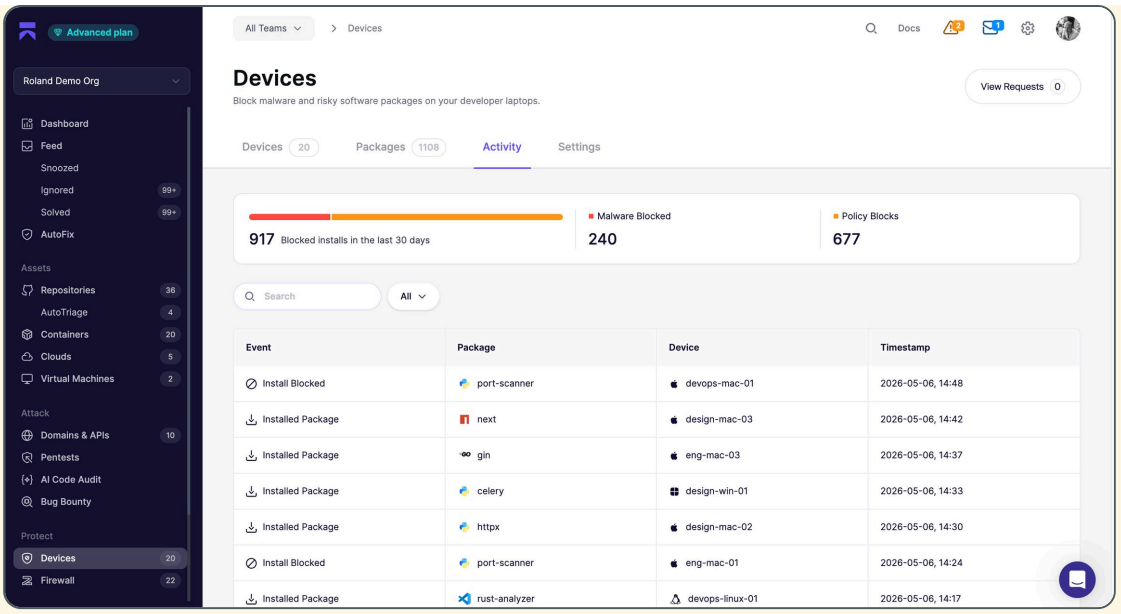
> 📍 In the Aikido console, open the original Android assessment, go to the overview tab, upload the rebuilt APK, then click 'Retest' on any finding from the initial pentest.

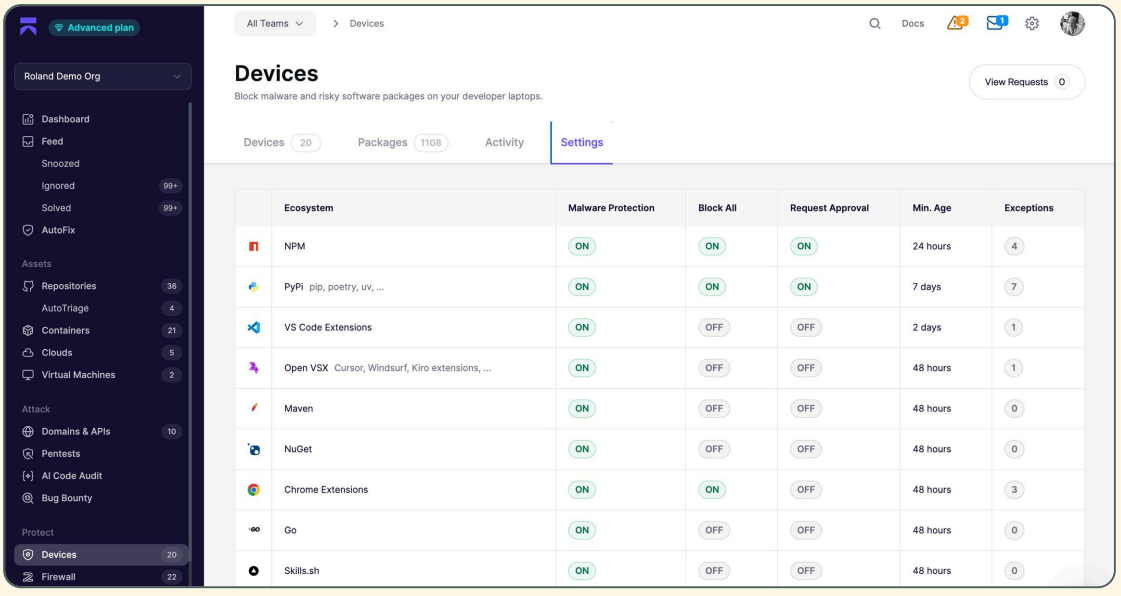
— Full mechanism, scope and a runnable retest workflow given. `launch-20260826-3c11bdb8`

---

02 **Device Protection agent for company devices** NEW 83

A new Device Protection agent provides real-time malware blocking across npm, PyPI, Maven, and NuGet package managers, powered by Aikido Intel, plus extension governance covering VS Code, OpenVSX, Chrome, and skills.sh extensions. A request-and-approval workflow pauses blocked installs and routes a one-click exception request to the admin Inbox inside Aikido, group-based policies let teams set different install and extension rules per team, and continuous device inventory gives an always-current view of every package and extension installed per device; the agent runs silently and only surfaces desktop notifications when a package is blocked or held for approval. Device Protection for Linux (beta) extends this to Linux desktops and servers, blocking malware at install time via `.deb` and `.rpm` packages.

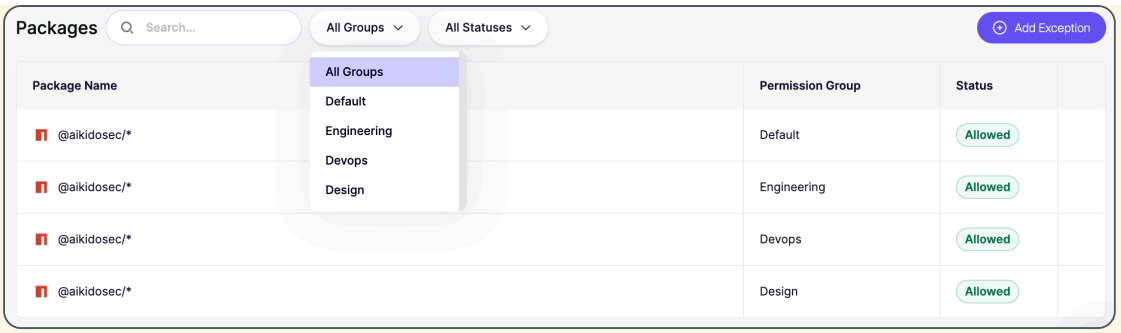




Inbox

Search | Pending | Handled | All Ecosystems

Ecosystem	Requested Package	Requested By	User Group	Actions
PyPi	requests	macos.shared		Declined
Open VSX	editor-code-assistant/eca	MacBook-Pro-bitterpanda63		Approved
PyPi	requests	eng-mac-03		Approved
PyPi	requests	eng-win-02		Approved
PyPi	requests	eng-mac-01		Declined
PyPi	test	eng-win-01		Declined



Review and act on a developer's blocked-install exception request from the admin Inbox.

> 📍 In the Aikido console, go to Inbox, find the pending install request, review the flagged package name and policy reason, then approve or deny the exception.

— Names mechanisms and package formats; includes an admin workflow example.

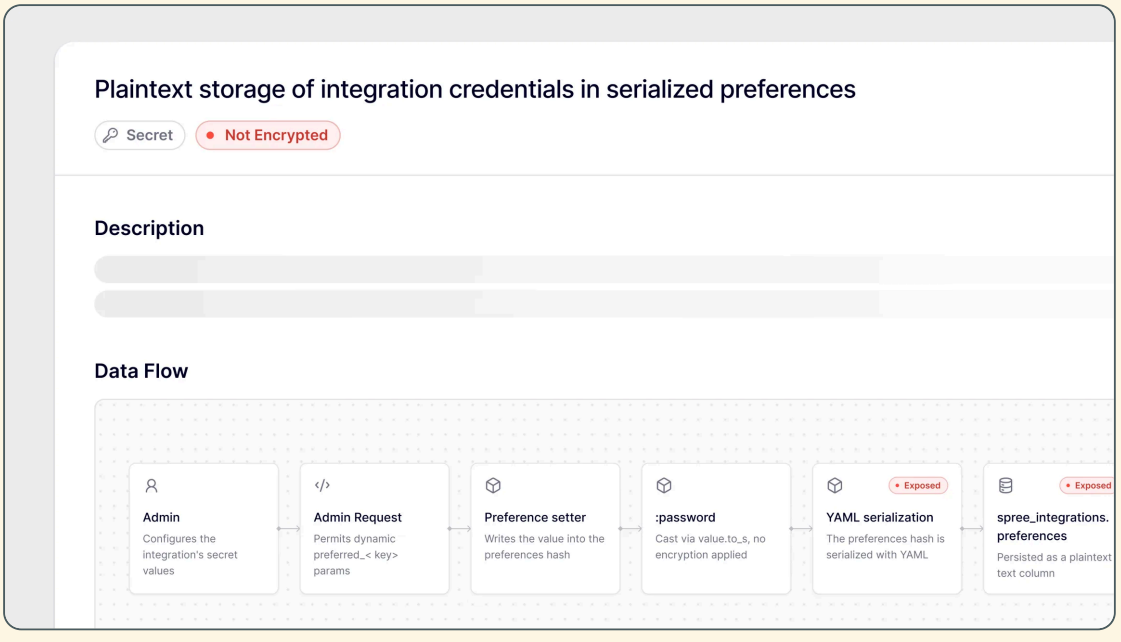
🚀 Protect Company Devices From Supply C… | 🚀 Protect Company Devices From Supply … | #54 - Device Protection for Linux and VM…

---

03 Code-based Data Exposure Audit (DSPM) NEW

83

Data Exposure Audit uses LLM agents to trace sensitive data from sources (request bodies, uploads, DB reads, config, app logic) into sinks (logs, third-party APIs, LLM prompts, caches, queues, storage), running entirely against already-connected repositories with no production data store access required. It detects credentials in URLs (tokens and API keys in query strings, redirects, and clone URLs), API over-exposure (endpoints returning more fields than the caller is entitled to, missing tenant ownership checks), unsafe logging and exports of unredacted PII and secrets, over-broad sharing of data with third parties and AI/LLM prompts and vector stores, weak storage (unencrypted or only base64/MD5/SHA-1 protected), incomplete deletion flows leaving downstream records, caches, and shadow copies intact, and over-permissioned IAM roles, OAuth scopes, and unauthenticated data stores. Each audit consumes credits from the Aikido Wallet.



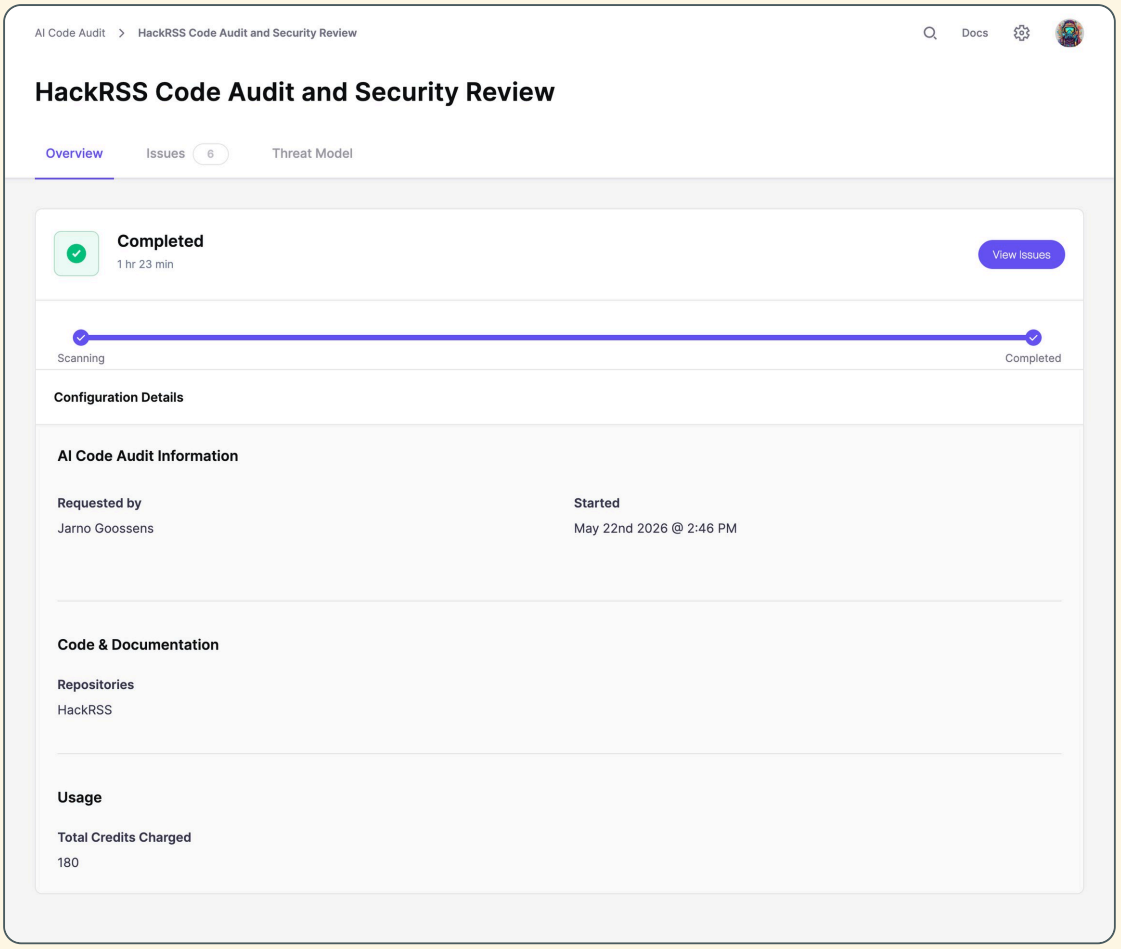
— Enumerates every detection category and the credit-consumption mechanism.

🚀 DSPM: Find data exposure risks withou… | 🚀 DSPM: Find data exposure risks witho…

04 Code Audit: pentest-grade source code reasoning NEW

76

Code Audit runs agentic, pentest-grade security reasoning directly on source code to find logic flaws (e.g. cross-tenant data leakage) that pattern-based scanners miss, with no staging URL, traffic replay, or agent deployment required. It surfaces code-based evidence, root cause analysis, and remediation guidance, and AutoFix can generate remediation PRs from findings in a few clicks. Pricing is shown upfront in credits before the audit starts, calculated from the size and complexity of the selected codebase, and it supports hard-to-test-by-URL targets including mobile apps, smart contracts, desktop apps, embedded software, and libraries.



Start a pentest-grade code audit on a repo with no live environment — useful for auditing a smart contract or mobile app codebase.

> 📍 In the Aikido dashboard, go to Code Audit, connect the target repository, review the upfront credit cost, then click Start Audit.

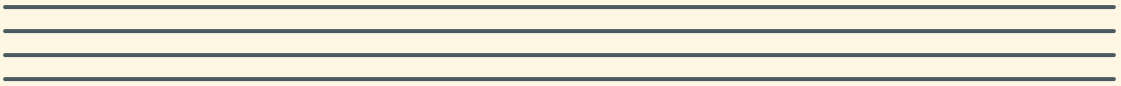
— Includes pricing mechanism, scope and a starting UI workflow.

🚀 Code Audit: Pentest-grade reasoning o… | 🚀 Code Audit: Pentest-grade reasoning …

05 **Secrets AutoTriage and CVE Exploitability Analysis agents** NEW

73

Secrets AutoTriage is an LLM agent that runs after entropy detection to assess whether a discovered secret is a real, exposed credential by evaluating liveness, scope, and surrounding code context, then applies defined rules or LLM reasoning to automatically ignore or escalate findings. CVE Exploitability Analysis is an LLM agent that evaluates whether a CVE is practically reachable and exploitable given how the codebase uses the affected package, and can automatically upgrade, downgrade, snooze, or ignore the finding; this assessment now also supports containers. Every agent run is logged with a short rationale explaining what was analyzed and what action was taken, so teams can review or override decisions.



— Explains both agents' reasoning and autonomous actions with audit logging.

🚀 Extreme Noise Reduction for Secrets a… 🚀 Extreme Noise Reduction for Secrets … #54 - Device Protection for Linux and VM…

THINNER COVERAGE BELOW

06 **Agentless GCP VM scanning (GA)** IMPROVED

54

GCP VM Scanning is now generally available, using an agentless setup that requires no SSH keys or agents, completing Aikido's VM scanning coverage across AWS, Azure, and GCP.

— Clear before/after (GA) but no exact setup steps given.

#54 - Device Protection for Linux and VM…

07 **Zen Firewall agent language support expanded** NEW

52

Zen Firewall agent releases are now available for Dotnet, Java, PHP, Ruby, Golang, and Node.

— Names every supported language but no further mechanism detail.

#54 - Device Protection for Linux and VM…

08 **Issue ignore workflows for PR/MR and license paths** NEW

48

Adds the ability to request ignore of an issue on a PR/MR with admin approval, and the ability to ignore license issues by path.

— Names two distinct ignore mechanisms but no exact UI/API path.

#54 - Device Protection for Linux and VM…

09 **OSCAL FedRAMP DAST report format** NEW

40

Adds an OSCAL FedRAMP DAST report format under Reports.

– Names the exact format and location but nothing further.

#54 - Device Protection for Linux and VM…

10 **Package Health support for NuGet and JetBrains** NEW 36

Package Health adds support for NuGet and JetBrains packages.

– Names two package ecosystems, no mechanism described.

#54 - Device Protection for Linux and VM…

11 **Bulk GitHub PR configuration API** NEW 30

New API to bulk update GitHub PR configuration across repositories.

– Names the API's purpose but no endpoint path or method.

#54 - Device Protection for Linux and VM…

12 **Hardened container image request API** NEW 24

New API to request a hardened container image.

– Bare mention of an API with no path or usage detail.

#54 - Device Protection for Linux and VM…

13 **New VSCode IDE extension release** IMPROVED 21

New VSCode IDE extension release.

– One-line mention with no detail on what changed.

#54 - Device Protection for Linux and VM…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 Zscaler

1 RELEASE · 2026-08-27

Secure Access (VPN / SASE / ZTNA)

Zscaler is a cloud-native security platform that provides secure internet access, zero trust networking, and threat protection for users, devices, and applications.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

Zscaler Zero Trust Browser can now discover, monitor, and block WebMCP AI-agent tool calls in-page, with inline policies matching on tool name, provider origin, and argument content — DLP-style inspection and enforcement of an agentic surface that was previously invisible.

> Zscaler added WebMCP Security Controls to its Zero Trust Browser, giving admins visibility into and inline control over AI agent tool calls on WebMCP-enabled sites.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 WebMCP inline policy enforcement NEW

95

Inline policies match on any of the seven `webmcp.*` attributes and apply one of three effects — Allow, Monitor, or Block — intercepting AI agent tool calls before the site's handler runs. Rules support matching on tool identity (`webmcp.tool.name`), argument content (`webmcp.tool.arguments`), provider origin (`webmcp.tool.provider_origin`), and cross-origin or untrusted-content signals (`webmcp.tool.is_cross_origin`, `webmcp.tool.has_untrusted_content_hint`) for DLP-style, argument-level inspection of agent payloads. Blocked calls return a structured failure to the agent ('The Zscaler Zero Trust Browser blocked this action for security reasons') with the exact rule reason (e.g., `webmcp.tool.name is equal to UPDATE_CART`) logged as a Low-severity Block detection.

Block a specific AI agent capability by tool name — for example, prevent any agent from invoking `update_cart` on any covered site, with the block surfaced inline to the user.

📍 In the Zscaler admin console, create a policy named 'Block WebMCP' with condition: `webmcp.tool.name` is equal to update_cart, effect: Block. The Detections view will show a Low-severity Block entry with Reason: `webmcp.tool.name` is equal to `UPDATE_CART` on every blocked invocation.

Apply argument-level (DLP-style) inspection to block agent tool calls whose JSON payload matches a keyword, while allowing the same tool with clean payloads.

📍 In the Zscaler admin console, create a policy rule with condition: `webmcp.tool.arguments` contains 'book', effect: Block. Calls to update_cart with payloads such as {"cart":{"line_items":[{"handle":"book-cat-graphic-tee-apple-

blossom-26f12103","quantity":1}]}} are blocked and the full argument payload is captured verbatim in the detection record.

— Full mechanism, named attributes/effects, and runnable policy examples.

2026-08-27-3c00b999

02 **WebMCP runtime monitoring in admin console** NEW 80

Every WebMCP tool invocation is captured with full context — the seven `webmcp.*` attributes plus user identity, device posture, extension version, browser, IP, geolocation, and URL — and surfaced in the Zscaler admin console as a searchable, filterable, exportable detection in the same event stream as existing platform telemetry.

Start in monitor mode to baseline all AI agent tool calls across sites before writing any block rules — every `webmcp.*` attribute is logged as an Informational/Allow detection without disrupting users.

📍 In the Zscaler admin console, create a policy named 'Monitor WebMCP' with all seven webmcp.* attributes enabled, default effect set to Allow, and assign it to the target user group. Open the Detections view to stream every WebMCP tool registration and invocation with full user, device, and argument context.

— Names console workflow and attributes, with a step-by-step monitor-mode example.

2026-08-27-3c00b999

03 **WebMCP capability discovery** NEW 68

The Zero Trust Browser Extension enumerates every tool a WebMCP-enabled site registers before any can run, capturing seven attributes per tool — `webmcp.tool.name`, `webmcp.tool.description`, `webmcp.tool.arguments`, `webmcp.tool_count`, `webmcp.tool.provider_origin`, `webmcp.tool.is_cross_origin`, and `webmcp.tool.has_untrusted_content_hint` — to build a structured inventory of the agentic surface per site and user.

— Names all seven attributes but no standalone usage example. 2026-08-27-3c00b999

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 JFrog

3 RELEASES · 2026-08-21 → 2026-08-28

SAST / SCA / Code Security

JFrog is a DevOps platform providing software supply chain solutions including artifact management, CI/CD integration, and security scanning.

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 5/5

JFrog Agent Package Resolution routes AI coding-agent dependency downloads (npm, PyPI, Maven, Go, Docker, Helm, NuGet) through Artifactory Curation automatically — including transitive and postinstall pulls — closing the supply-chain blind spot that Claude Code and Cursor sessions reopened, with full audit trails.

> JFrog opened a new Preview early-access track for upcoming capabilities, shipped Agent Package Resolution to route AI coding agents' dependency downloads through Artifactory governance, and added On-Behalf-Of token exchange so Amazon Bedrock AgentCore Gateway can propagate per-user identity into Artifactory.

└──▷ WHAT SHIPPED · 3 FEATURES most completely described first ? what's the number?

01 Agent Package Resolution for AI coding agents NEW 95

Agent Package Resolution (preview), part of the JFrog Agent Plugin, routes AI coding agent dependency downloads (Claude Code, Cursor) through JFrog Artifactory instead of public registries, covering npm, PyPI, Maven, Go, Docker Hub, Helm, and NuGet. Session steering automatically supplies the Artifactory URL for each governed package type at the start of every agent session with no per-project configuration, while the new `jf setup` command persistently binds the local package manager to an Artifactory repository for a given project across sessions and tools. JFrog Curation policy enforcement is extended server-side to all agent-driven downloads — blocking malicious packages, unapproved licenses, and policy-violating packages regardless of whether session steering or `jf setup` is active — and covers indirect/transitive installs such as postinstall scripts, `pip install -r requirements.txt`, and `npm ci`. Every agent-driven download, block, and policy decision is captured in audit trails alongside existing human-driven build logs.

One-time project setup: bind your local package managers to Artifactory so that every subsequent agent session — and any indirect installs — resolves through the governed repository without manual configuration.

```
$ jf setup
```

Enable Agent Package Resolution for your team via the JFrog Agent Plugin on GitHub so Claude Code or Cursor sessions automatically receive the correct Artifactory URLs at

startup.

> 📍 In the JFrog Platform, install the JFrog Agent Plugin from GitHub, then under Agent Package Resolution settings choose which package types (npm, PyPI, Maven, Go, Docker, Helm, NuGet) are governed, and enable zero-touch setup to push Artifactory URLs to agent sessions automatically.

— Names exact command, config surfaces and package types covered.

Coding Agents Just Reopened Your Softwar…

---

02 **On-Behalf-Of identity propagation with Bedrock AgentCore Gateway**

NEW 95

Artifactory now supports On-Behalf-Of (OBO) OAuth 2.0 token exchange (RFC 8693) between Amazon Bedrock AgentCore Gateway and JFrog Artifactory Access, so AI agent requests reach Artifactory authenticated as the originating user rather than a shared service account, preserving audit fidelity, least privilege, and containment. AgentCore Gateway registers Artifactory's MCP server as a target configured with `onBehalfOfTokenExchangeConfig`, exchanging each inbound user JWT for a short-lived per-user JFrog token; the `listingMode` field (requiring a recent AWS CLI) enables dynamic tool listing so discovery runs under the user's token instead of a shared credential. JFrog Access OIDC identity mappings integrate with AgentCore Identity by trusting the external IdP issuer (`$IDP_ISSUER`) and mapping verified access-token claims (e.g. `preferred_username` , `team`) to JFrog users, groups, or project roles via the JFrog OIDC integration API, requiring the token's `aud` claim to equal the AgentCore Gateway URL (`$GATEWAY_URL`) and `iss` to match `$IDP_ISSUER` . This approach is compared against header whitelisting and three-legged OAuth, with OBO being the only option delivering verifiable per-user identity without a second login; it requires JFrog EnterpriseX or higher.

— Names every config field and claim but no runnable command.

Propagating User Identity From AI Agents…

---

03 **Preview early-access program** NEW 72

JFrog launched a 'Preview' early-access track offered as Closed Preview (invite-only) or Open Preview (all eligible customers), both free to opt into. Customers activate it in the MyJFrog portal under `Settings > Preview` by enabling the Enable Preview toggle, accepting the Preview Terms of Service in the 'Enable JFrog Preview' dialog, and then selecting specific preview features to activate for their subscription.

Opt in to JFrog Preview to gain immediate access to newly released capabilities before they reach general availability.

📍 In MyJFrog, go to Settings › Preview. Click the 'Enable Preview' toggle. In the 'Enable JFrog Preview' dialog box, open and read the Preview Terms of Service. Select the checkbox to accept the terms and conditions, and click 'Agree and Enable'. Then select specific preview features to activate in your subscription.

— Exact portal path and toggle given, but scope of features unnamed.

launch-20260828-b987b724

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 defectdojo

3 RELEASES · 2026-08-24 → 2026-08-27

Vulnerability Management

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

DefectDojo Pro's Sensei adds CSPM: bulk-onboard AWS, Azure, and GCP accounts for cloud misconfiguration scanning with 'Fix in Cloud' direct remediation through the provider API — gated by a human-approval dialog and a revert button — alongside a new EPSS/KEV-enriched Vulnerability Explorer for one-row-per-CVE prioritization.

> DefectDojo Pro launched Sensei CSPM for cloud misconfiguration scanning and one-click remediation across AWS, Azure and GCP, alongside a new Vulnerability Explorer with EPSS/KEV enrichment and a beta Root Cause Correlation clustering feature; the core platform also tightened endpoint permissions, added MFA disable verification, and expanded its documentation.

└──▷ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 **Sensei CSPM cloud scanning and remediation** NEW 90

DefectDojo Pro adds Cloud Connection onboarding to connect AWS Organizations, Azure subscriptions, and GCP projects in bulk using a shared read-only credential that auto-discovers and onboards all visible accounts at once. On-demand cloud posture scanning from the Sensei hub produces deduplicated, prioritized findings tied to a specific cloud resource (S3 bucket, security group, etc.) and updated in place on re-scan. A new 'Fix in Cloud' action applies changes directly through the provider API using a separate write credential (independent of the read-only scan credential), gated by a mandatory human-approval dialog that shows the exact action, target resource, and required cloud permissions before execution; the existing IaC pull-request remediation path is retained for resources managed as code, so Sensei opens a PR instead of touching the live resource when source exists.

Onboard all AWS accounts visible to an Organizations management credential and run an initial posture scan from the Sensei hub.

> 📍 In DefectDojo Pro, go to Sensei › Cloud Connections › Add Connection, select AWS, enter a read-only Organizations credential, and choose 'Discover and onboard all member accounts'. Then navigate to Sensei › Hub and select 'Run Scan' against the newly onboarded accounts.

Apply and later revert a Fix in Cloud action (e.g., blocking public S3 access) with human approval and a full audit trail.

📍 In DefectDojo Pro, open a cloud posture finding for an S3 bucket, click 'Fix in Cloud', review the approval dialog (action: block public access, resource: the specific bucket, required permissions listed), approve, then navigate to Sensei › Cloud Remediations to see the ledger entry and use the per-row 'Revert' button if needed.

— Full mechanism and named credential/approval flow, minor gaps on scan cadence limits

`launch-20260827-8a253852`

---

02 **Vulnerability Explorer with EPSS/KEV enrichment** NEW

86

DefectDojo Pro adds the Vulnerability Explorer, a deduplicated, one-row-per-vulnerability-ID view showing EPSS Score, EPSS Percentile, Known Exploited, Ransomware Used, and KEV Date columns plus a Total Findings count that links directly to the filtered findings list. A `Sync KEV/EPSS data` button pulls the latest CISA KEV and FIRST EPSS data on demand in addition to an automatic daily sync. SLA configuration now supports Cap by KEV due date, tightening a KEV-listed finding's deadline to CISA's official due date, and the view distinguishes unenriched vulnerabilities (blank column) from confirmed non-exploited ones (red ✕ under Known Exploited).

Quickly surface the highest-exploitation-probability CVEs across your entire portfolio for same-day triage.

📍 In DefectDojo Pro, open the Vulnerability Explorer, click the 'EPSS Score' column header to sort descending, then click 'Total Findings' on the top row to jump directly to the filtered findings list for that CVE.

Force an immediate refresh of KEV and EPSS data after CISA publishes a new catalog update, without waiting for the nightly sync.

📍 In DefectDojo Pro, open the Vulnerability Explorer and click the 'Sync KEV/EPSS data' button at the top of the page.

— Named columns, button, and SLA config key with clear steps to use them

`launch-20260825-d1ed3339`

---

03 **Root Cause Correlation (Beta)** NEW

67

DefectDojo Pro adds a Beta Root Cause Correlation feature that groups findings sharing the same component-at-version, CVE, infrastructure resource, or weakness class at the same URL into ranked clusters showing active findings, affected assets, and aggregate risk. It is enabled by a superuser under Settings > Feature Flags.

Enable Root Cause Correlation (Beta) so that clusters of findings sharing the same vulnerable component are grouped and prioritized together.

> 📍 In DefectDojo Pro, navigate to Settings › Feature Flags and enable Root Cause Correlation (requires superuser account).

— Clear enablement path but clustering algorithm details are thin `launch-20260825-d1ed3339`

THINNER COVERAGE BELOW

04 Expanded documentation across product areas IMPROVED

59

DefectDojo's documentation was expanded to cover typed Custom Fields and their integration with Rules Engine 2.0, multi-Asset finding fan-out in Smart Upload workflows, AI Model Settings and supported LLM providers in Sensei, GMP over SSH support for OpenVAS/Greenbone connectors, customizable page layouts and Form Configuration in navigation, the Quick Export dialog for named reports, the rebuilt Asset hierarchy page, and the four-permission rule tier in Rules Engine 2.0. It also adds an RFC 9116 `security.txt` file to `docs.defectdojo.com`, adds Brazilian Portuguese, Simplified Chinese, and Italian translations of the core guides, and splits connector documentation into one page per connector.

— Documentation-only changes; names topics but no functional mechanism `3.2.300`

05 Endpoint delete/rename scoped to product IMPROVED

33

Endpoint delete and rename operations are now scoped to the requesting product, restricting cross-product mutations that were previously possible.

— Behavior change stated but no endpoint name or migration detail given `3.2.300`

06 MFA disable requires verification code IMPROVED

31

Disabling MFA now requires entering a verification code, adding a confirmation step intended to protect account security.

— Clear behavior change but no UI path or code details specified `3.2.300`

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · August 28, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe