THE CYBER TOOLCHAIN

NO. 037

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 6, 2026 · EVERY DAY

The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 38 of 41 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06

NOTES ↗

Horizon3 NodeZero is an autonomous security tool that continuously identifies and exploits vulnerabilities in networks to assess organizational risk.

> Horizon3 NodeZero's biggest addition this window is Tripwires — automatically deployed decoy files and credentials that flag intruder interaction during pentests — alongside expanded Rapid Response exploit coverage, Active Directory detection improvements, and MCP Server, MSP, and Vulnerability Management Hub upgrades.

└──▷ WHAT SHIPPED · 7 FEATURES most completely described first ? what's the number?

01 NodeZero Tripwires decoy files and credentials NEW

75

NodeZero automatically deploys decoy files and credentials — up to two tripwires per asset — along critical attack paths during autonomous pentests, placed at points where NodeZero has gained access via a Remote Access Tool (RAT) or direct file write; any interaction triggers an immediate alert. Tripwires are included in NodeZero Elite and Pro license packages and available as an add-on for other tiers. A subsequent update introduced more-flexible Tripwire deployment that reduces reliance on the RAT, cutting Endpoint Detection and Response (EDR) noise during deployments.

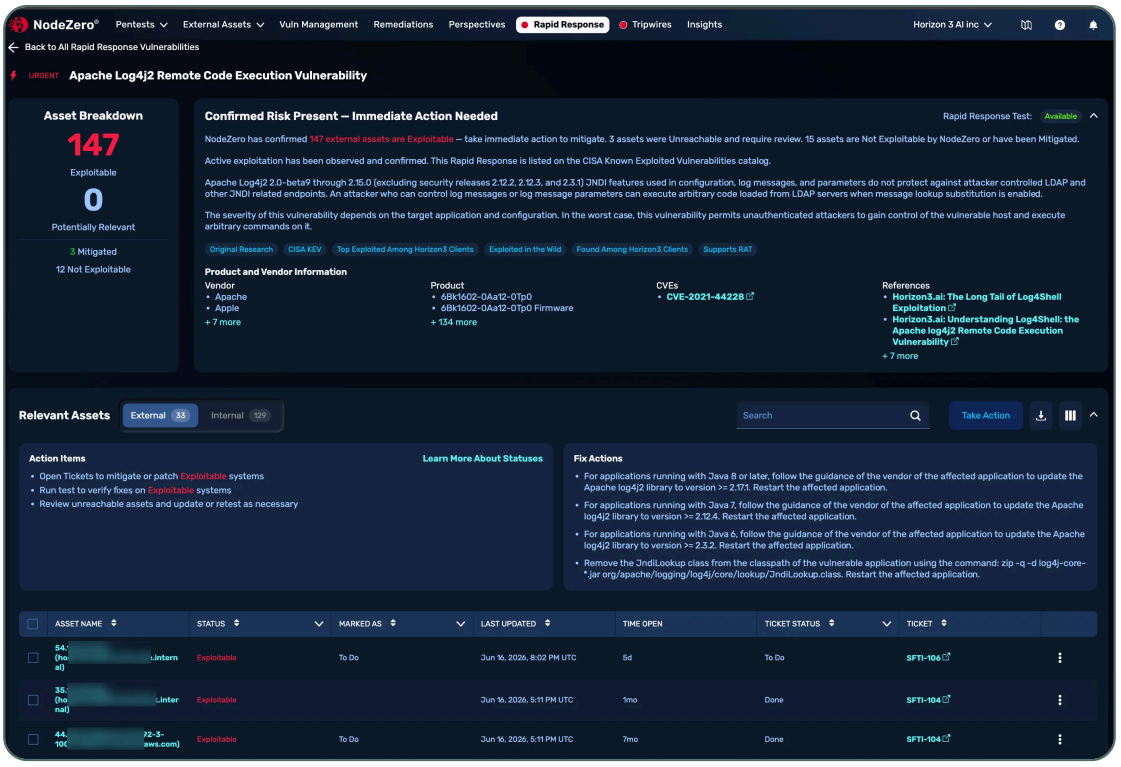
— Names mechanism, per-asset limit, license tiers; no config surface given.

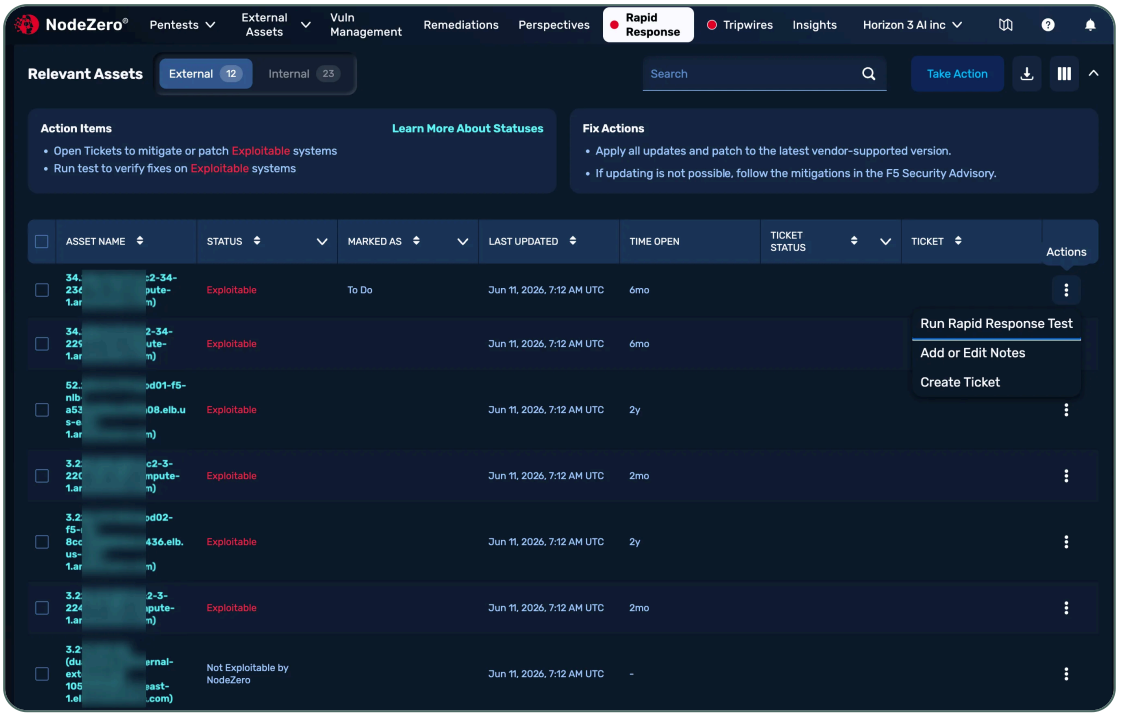
snapshot-20260806

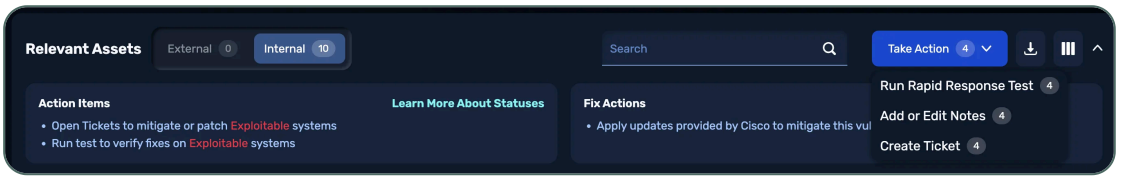
02 Expanded Rapid Response exploit coverage NEW

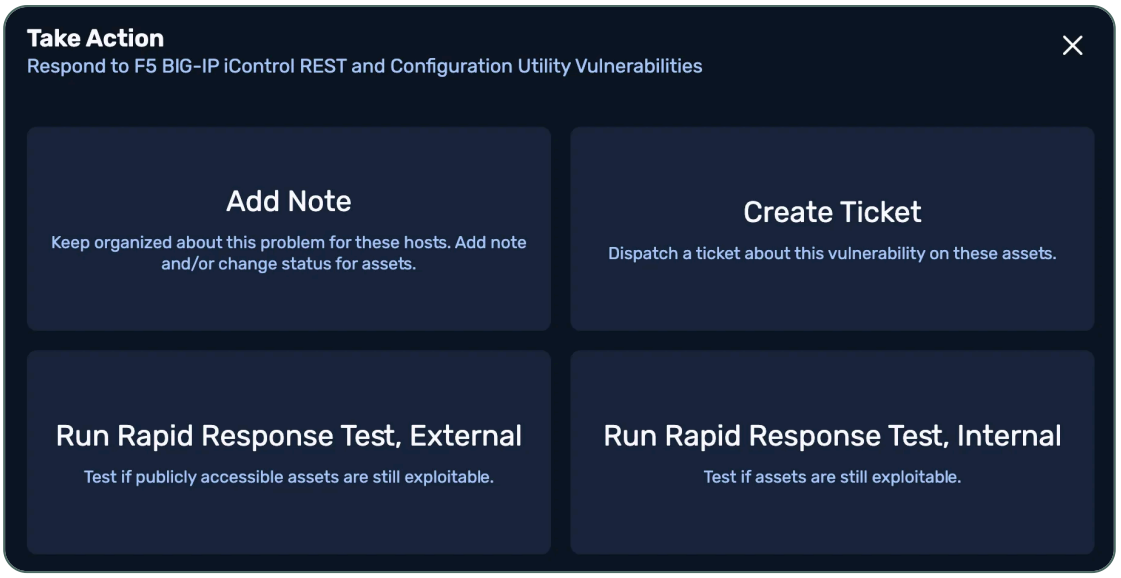
65

Rapid Response exploitability testing now covers Cisco Unified Communications Manager, SonicWall SMA1000, Adobe ColdFusion, and ManageEngine ADAudit Plus, viewable via Rapid Response Details pages showing Asset Breakdown, CVE or H3 IDs, References, and a Relevant Assets table with Action Items and Fix Actions guidance, plus a Take Action menu to run tests, add notes, or create tickets.


NodeZero
Pentests
External Assets
Vuln Management
Remediations
Perspectives
Rapid Response
Tripwires
Insights
Horizon 3 AI inc
Back to All Rapid Response Vulnerabilities
URGENT
Apache Log4j2 Remote Code Execution Vulnerability
Asset Breakdown
147
Exploitable
0
Potentially Relevant
3 Mitigated
12 Not Exploitable
Confirmed Risk Present — Immediate Action Needed
Rapid Response Test: Available
NodeZero has confirmed 147 external assets are Exploitable — take immediate action to mitigate. 3 assets were Unreachable and require review. 15 assets are Not Exploitable by NodeZero or have been Mitigated.
Active exploitation has been observed and confirmed. This Rapid Response is listed on the CISA Known Exploited Vulnerabilities catalog.
Apache Log4j2 2.0-beta9 through 2.15.0 (excluding security releases 2.12.2, 2.12.3, and 2.3.1) JNDI features used in configuration, log messages, and parameters do not protect against attacker controlled LDAP and other JNDI related endpoints. An attacker who can control log messages or log message parameters can execute arbitrary code loaded from LDAP servers when message lookup substitution is enabled.
The severity of this vulnerability depends on the target application and configuration. In the worst case, this vulnerability permits unauthenticated attackers to gain control of the vulnerable host and execute arbitrary commands on it.
Original Research
CISA KEV
Top Exploited Among Horizon3 Clients
Exploited in the Wild
Found Among Horizon3 Clients
Supports RAT
Product and Vendor Information
Vendor
Apache
Apple
+ 7 more
Product
6Bk1602-0Aa12-0Tp0
6Bk1602-0Aa12-0Tp0 Firmware
+ 134 more
CVEs
CVE-2021-44228
References
Horizon3.ai: The Long Tail of Log4Shell Exploitation
Horizon3.ai: Understanding Log4Shell: the Apache log4j2 Remote Code Execution Vulnerability
+ 7 more
Relevant Assets
External 33
Internal 129
Search
Take Action
Action Items
Learn More About Statuses
Open Tickets to mitigate or patch Exploitable systems
Run test to verify fixes on Exploitable systems
Review unreachable assets and update or retest as necessary
Fix Actions
For applications running with Java 8 or later, follow the guidance of the vendor of the affected application to update the Apache log4j2 library to version >= 2.17.1. Restart the affected application.
For applications running with Java 7, follow the guidance of the vendor of the affected application to update the Apache log4j2 library to version >= 2.12.4. Restart the affected application.
For applications running with Java 6, follow the guidance of the vendor of the affected application to update the Apache log4j2 library to version >= 2.3.2. Restart the affected application.
Remove the JndiLookup class from the classpath of the vulnerable application using the command: zip -q -d log4j-core-*.jar org/apache/logging/log4j/core/lookup/JndiLookup.class. Restart the affected application.
ASSET NAME
STATUS
MARKED AS
LAST UPDATED
TIME OPEN
TICKET STATUS
TICKET
Exploitable
To Do
Jun 16, 2026, 8:02 PM UTC
5d
SFTI-106
Exploitable
Jun 16, 2026, 5:11 PM UTC
1mo
Done
SFTI-104
Exploitable
To Do
Jun 16, 2026, 5:11 PM UTC
7mo
Done
SFTI-104


NodeZero
Pentests
External Assets
Vuln Management
Remediations
Perspectives
Rapid Response
Tripwires
Insights
Horizon 3 AI inc
Relevant Assets
External 12
Internal 23
Search
Take Action
Action Items
Learn More About Statuses
Open Tickets to mitigate or patch Exploitable systems
Run test to verify fixes on Exploitable systems
Fix Actions
Apply all updates and patch to the latest vendor-supported version.
If updating is not possible, follow the mitigations in the F5 Security Advisory.
ASSET NAME
STATUS
MARKED AS
LAST UPDATED
TIME OPEN
TICKET STATUS
TICKET
Actions
Exploitable
To Do
Jun 11, 2026, 7:12 AM UTC
6mo
Run Rapid Response Test
Add or Edit Notes
Create Ticket
Exploitable
Jun 11, 2026, 7:12 AM UTC
6mo
Exploitable
Jun 11, 2026, 7:12 AM UTC
2y
Exploitable
Jun 11, 2026, 7:12 AM UTC
2mo
Exploitable
Jun 11, 2026, 7:12 AM UTC
2y
Exploitable
Jun 11, 2026, 7:12 AM UTC
2mo
Not Exploitable by NodeZero
Jun 11, 2026, 7:12 AM UTC
-


Relevant Assets
External 0
Internal 10
Search
Take Action 4
Run Rapid Response Test 4
Add or Edit Notes 4
Create Ticket 4
Action Items
Learn More About Statuses
Open Tickets to mitigate or patch Exploitable systems
Run test to verify fixes on Exploitable systems
Fix Actions
Apply updates provided by Cisco to mitigate this vul



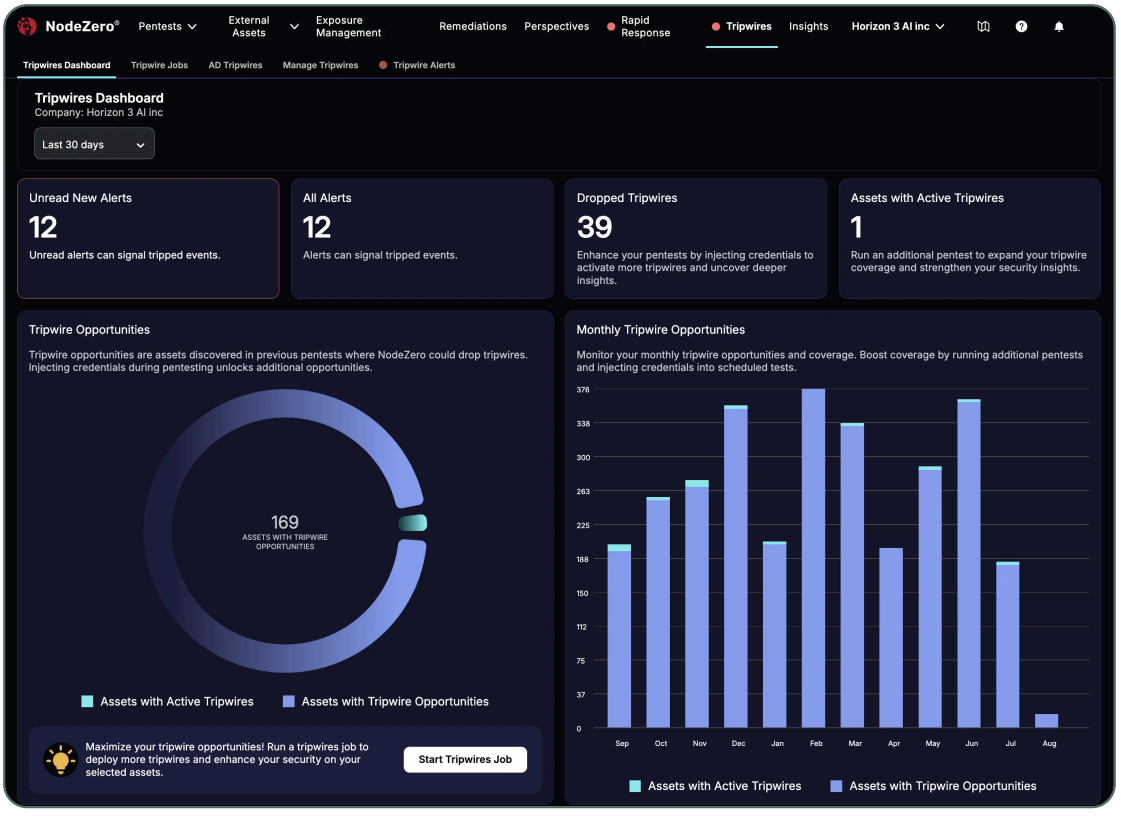
— Names covered products and UI actions, no API detail. `snapshot-20260806`

THINNER COVERAGE BELOW

03 Tripwires dashboard in NodeZero Portal NEW

45

A Tripwires dashboard in the NodeZero Portal, opened from the portal's top navigation, gives visibility into deployed tripwires across the environment.



— Names exact UI path but no further behaviour detail. `product docs`

04 Vulnerability Management Hub port info and sorting IMPROVED

40

The Vulnerability Management Hub adds port information and improved sorting for completed pentests.

— Names the hub and new fields but no navigation steps. snapshot-20260806

05 **Active Directory group and DCSync detection improvements**

IMPROVED

35

Active Directory enhancements improve group enumeration, NTLM posture assessment, and DCSync detection.

— Names three capabilities but no mechanism or scope detail. snapshot-20260806

06 **MSP contract, asset, and co-branding controls** IMPROVED

35

MSP enhancements strengthen contract enforcement, asset allocation, and co-branding controls for managed service provider environments.

— Names three MSP controls without configuration or UI detail. snapshot-20260806

07 **MCP Server authentication and toolset updates** IMPROVED

30

MCP Server updates deliver stronger authentication, a streamlined toolset, and expanded documentation for safer integrations.

— Vague on what auth or toolset changes actually entail. snapshot-20260806

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

ProjectDiscovery subfinder

Feedback ⓘ

1 RELEASE · 2026-08-05 NOTES ↗

Fast passive subdomain enumeration tool.

> Subfinder v2.15.0 adds a new passive discovery source and machine-readable output for listing supported sources.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 **Machine-readable source listing via -ls -oJ** NEW 75

Running `subfinder -ls -oJ` outputs the available passive sources in JSONL format, useful for feeding into automation or auditing which sources subfinder supports.

Get a machine-readable list of all available passive sources to feed into automation or audit which sources subfinder supports.

```
$ subfinder -ls -oJ
```

— Exact flag and output format given, with runnable example. v2.15.0

02 **New shodanct passive source** NEW 75

Adds `shodanct` as a new passive subdomain discovery source, leveraging Shodan's certificate transparency data; can be invoked with `subfinder -d example.com -s shodanct`.

Run subdomain enumeration using only the new shodanct source to leverage Shodan's certificate transparency data.

```
$ subfinder -d example.com -s shodanct
```

— Named source and exact command shown, mechanism briefly noted. v2.15.0

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Snyk

Feedback ⓘ 1 RELEASE · 2026-08-06 DOCS ↗ BLOG ↗

Snyk is a developer security platform that finds and fixes vulnerabilities in open source dependencies and container images.

> Snyk added NTLM v2 authentication support for Web targets in Snyk API & Web, allowing authenticated DAST scans of Windows-protected internal applications.

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

01 **NTLM v2 authentication for Web targets** NEW 83

Snyk API & Web now supports NTLM v2 authentication for Web targets, configurable via `Target Settings > Authentication` tab in the `Basic Auth / NTLM` module, with fields for username, password, domain, and workstation. This enables authenticated scans of Windows-protected internal applications without separate tooling.

Enable authenticated scanning of a Windows NTLM-protected internal application that was previously unscanned.

> 📍 In the Snyk web UI, go to Target Settings › Authentication tab, locate the Basic Auth / NTLM module, select the NTLM option, and enter the username, password, domain, and workstation credentials for the target application.

— Names exact UI path and auth fields but no API/flag. 2026-08-06-a1520ef7

WAS THIS USEFUL? 👍 👎

Portkey

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06 NOTES ↗

Portkey is an AI gateway providing observability, security, and reliability features for LLM applications with multi-model support.

> Portkey's March 2026 release adds vault-backed Secret References, weekly and endpoint-scoped rate limits, cross-cloud GCS logging via GCP Workload Identity Federation, new Zscaler and Akto guardrail integrations, and expanded model/provider support including DeepInfra, DeepSeek, Bedrock, Vertex AI, and Azure AI Foundry rerank.

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

01 **Weekly rate-limit budgets** NEW 45

Adds weekly (`rpw`) budget and usage windows so rate-limit caps align with weekly spend planning, not only per-minute or monthly aggregates.

— Names the `rpw` unit but no config path or command. snapshot-20260806

02 **Analytics for archived workspaces and slug filters** IMPROVED 45

Enables analytics for archived workspaces and accepts workspace slugs (alongside IDs) in saved filters, keeping reporting stable as workspaces are wound down or renamed.

— Describes behaviour but no filter syntax or API path. snapshot-20260806

03 **Removal of forwarded header to prevent loops** BREAKING 45

The gateway no longer forwards the `x-portkey-forward-headxers` header, preventing header-forwarding loops; any setup relying on that header being passed through will stop working.

— Names the removed header but no migration path given. snapshot-20260806

└──▷ BREAKING ON UPGRADE

! The gateway no longer forwards the `x-portkey-forward-headxers` header, preventing header-forwarding loops — any setup relying on that header being passed through will stop working.

WAS THIS USEFUL? 👍 👎

Openlayer

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06 NOTES ↗

Openlayer is a platform for monitoring and evaluating large language model performance, quality, and safety in production.

> Openlayer shipped a Gantt-style trace timeline view that visualizes span timing and nesting to help spot slow steps and unexpected branches.

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

01 **Gantt-style trace timeline view** NEW 45

A new trace timeline view lines up every span by start time, duration, and nesting so slow steps and unexpected branches are immediately visible at a glance.

— Describes behavior but no navigation path or config named. snapshot-20260806

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ 2 RELEASES · 2026-08-05 → 2026-08-06 BLOG ↗

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities.

> Veracode expanded its EASM module with real-time scan notifications, AI-powered risk prioritization, scan scheduling, and a new DAST Candidate export path into DAST Essentials, while its training platform shipped Get Lesson Progress API enhancements, OWASP Top 10:2025 content, and deeper Security Labs integration into the Veracode Platform via IAM.

└──▷ **WHAT SHIPPED · 25 FEATURES** most completely described first ? what's the number?

01 **Get Lesson Progress API enhancements** IMPROVED 95

Adds `updatedAfter` parameter to support incremental loads of lesson progress data and `excludeNotStarted=true` to filter out lessons that have not been started. The response now also includes `stack` and `updated_on` fields, and surfaces `Not Started` status when both `start_time` and `end_time` are null.

— Names exact params, fields and status logic. Training updates (2026-08-06)

02 **Security Labs default URL change** BREAKING 80

The default Security Labs URL changed from `https://securitylabs.veracode.com/` to `https://securitylabs.veracode.com/security-labs/`; customers using SSO with a stand-alone account must update their configured URL.

— Exact old/new URLs plus migration note given. Training updates (2026-08-06)

03 **DAST Candidate export to DAST Essentials** NEW 70

Enables export of web applications from EASM directly into DAST Essentials by marking assets as DAST Candidates.

— Names target product and exact marking mechanism. EASM updates (2026-08-05)

04 **OWASP Top 10:2025 Security Labs lessons** NEW 60

New Security Labs lessons cover all ten OWASP Top 10:2025 categories, including two new categories: A03 (Software Supply Chain Failures) and A10 (Mishandling of Exceptional Conditions).

— Names categories but only prose description. Training updates (2026-08-06)

05 **OWASP Top 10:2025 eLearning courses** NEW 60

New eLearning courses aligned to OWASP Top 10:2025 are available in English, Spanish, French, Portuguese, and Chinese, covering Software Security Awareness and secure

coding for .NET, Java, JavaScript, PHP, and Python.

— Lists languages and topics, no further mechanism. Training updates (2026-08-06)

06 **eLearning skills assessment quizzes** NEW 60

New assessment quiz courses (skills assessments) are available in English, Portuguese, and Spanish for Static Analysis, Dynamic Analysis, SCA agent-based scanning, IDE scanning, and user/application setup.

— Named topics and languages, no navigation detail. Training updates (2026-08-06)

07 **CWE tagging on Security Labs lessons** NEW 60

Adds CWE tagging to Security Labs lessons, enabling admins to search lessons by CWE and learners to see which CWEs a lesson addresses.

— Explains search and display behavior for CWEs. Training updates (2026-08-06)

08 **OWASP API Security Top 10 upgraded to 2023** IMPROVED 60

Upgrades OWASP API Security Top 10 coverage in Security Labs from the 2019 list to the 2023 list, adding lessons for three new categories.

— Named version change and category count. Training updates (2026-08-06)

THINNER COVERAGE BELOW

09 **Security Labs terminology rename** BREAKING 55

In Security Labs, Teams are now called Organizations and Roles are now called Teams; any documentation, automation, or integrations referencing these terms must be updated.

— Exact term mapping given, no further mechanism. Training updates (2026-08-06)

10 **Security Labs user/team management via IAM** NEW 55

Integrates Security Labs user and team management into the Veracode Platform via Veracode IAM, accessible under Security Training after migration.

— Names IAM and Security Training location, little else. Training updates (2026-08-06)

11 **33 deprecated eLearning courses removed** DEPRECATED 55

33 deprecated eLearning courses (OWASP 2017, OWASP Top 10 2021, and language-specific Secure Coding courses for .NET, Java, JavaScript, PHP, and Python) are no longer available as of June 10, 2026. OWASP Top 10 (2021) Security Labs lessons are also no longer available as of that date.

— Specific count and date, no replacement steps. Training updates (2026-08-06)

12 **In-app guidance for campaigns and reporting** NEW 50

New in-app guidance in Security Labs for creating campaigns and using reporting; requires third-party cookies enabled in Chrome, Edge, Firefox, Safari, and Internet Explorer 9+.

— Lists browser requirement, no exact UI path. Training updates (2026-08-06)

13 **Dashboard Heatmap, World map, Things widgets** NEW 50

Adds Dashboard Heatmap, World map, and Things widgets with filterable data views and a shortcut to create reports.

— Named widgets, limited on filter/report detail. EASM updates (2026-08-05)

14 **Container Security in Docker lesson series** NEW 48

New Security Labs lesson series on Container Security in Docker, covering capabilities, users, virtualization, networking, resources, and inventory across dozens of hands-on lessons.

— Topic list given, no lesson count or path. Training updates (2026-08-06)

15 **Real-time in-app scan notifications** NEW 45

Adds real-time in-app notifications when scans complete, fail, or are cancelled.

— States trigger events, no further detail. EASM updates (2026-08-05)

16 **Findings view with prioritization workflows** NEW 45

Adds a Findings view with improved prioritization workflows and expanded reporting capabilities.

— Names the view, but workflow details are thin. EASM updates (2026-08-05)

17 **Business units management** NEW 45

Adds business units management: create, assign, and manage business units with domain associations.

— Names CRUD actions and domain association. EASM updates (2026-08-05)

18 **Deep Discovery work queue improvements** IMPROVED 43

Enhances Deep Discovery with an improved work queue, selection toggles, and clearer result indicators.

— Named UI elements, no workflow specifics. EASM updates (2026-08-05)

19 **Security Labs team search for admins** NEW 35

Adds Security Labs team search capability for Security Labs administrators and managers.

— Brief description, no mechanism or scope. Training updates (2026-08-06)

20 **Scan scheduling for orchestration** NEW 33

Adds scan scheduling to streamline scan orchestration across projects.

— Names capability but no config or frequency detail. EASM updates (2026-08-05)

21 **Issues capability for remediation** NEW 33

Adds Issues capability to streamline remediation efforts and improve cyber hygiene focus.

— Names the feature, mechanism unspecified. EASM updates (2026-08-05)

22 **Security Labs community edition deprecated** DEPRECATED 30

The Security Labs community edition will be deprecated on July 31, 2026.

— Date given but no migration guidance. Training updates (2026-08-06)

23 **Code syntax highlighting in lessons** IMPROVED 30

Adds code syntax highlighting to Security Labs lessons, tailored per programming language.

— Minimal description of a UI-only change. Training updates (2026-08-06)

24 **Multi-project visibility for risk management** NEW 30

Adds multi-project visibility for centralized risk management across all projects.

— Generic description, no mechanism named. EASM updates (2026-08-05)

25 **AI-powered intelligence for risk prioritization** NEW 20

Adds AI-powered intelligence for accelerated risk prioritization.

— Vague marketing-level description of mechanism. EASM updates (2026-08-05)

└──▷ BREAKING ON UPGRADE

! The default Security Labs URL has changed from `https://securitylabs.veracode.com/` to `https://securitylabs.veracode.com/security-labs/`; customers using SSO with a stand-alone account must update their configured URL.

! In Security Labs, Teams are now called Organizations and Roles are now called Teams; any documentation, automation, or integrations referencing these terms must be updated.

! The Security Labs community edition will be deprecated on July 31, 2026.

! 33 deprecated eLearning courses (OWASP 2017, OWASP Top 10 2021, and language-specific Secure Coding courses for .NET, Java, JavaScript, PHP, and Python) are no longer available as of June 10, 2026.

! OWASP Top 10 (2021) Security Labs lessons are no longer available as of June 10, 2026.

WAS THIS USEFUL?  

◆ Web Application Security

ProjectDiscovery Katana

Feedback ⓘ 1 RELEASE · 2026-08-05

NOTES ↗ CODE ↗

A next-generation crawling and spidering framework.

> Katana v1.7.0 expands its knowledge base with secrets and endpoint extractors, adds content similarity deduplication to cut near-duplicate pages, and opens up headless engine callbacks with a Context field for external lifecycle control.

└──▷ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Content similarity deduplication of crawled pages NEW

100

The `-pcs` / `-page-content-similar` flag (also aliased as `-sdd` / `-similarity-deduplication`) enables Layer-2 content similarity filtering to deduplicate near-identical pages during crawls. Mode is set via `-pcsm` / `-page-content-similar-mode` (simhash, tfidf, or bm25; default `simhash`), tuned with `-pcsd` / `-page-content-similar-distance` (SimHash max Hamming distance, default 3), `-pcst` / `-page-content-similar-threshold` (TF-IDF/BM25 min score 0–1, default 0.85), and `-pcsn` / `-page-content-similar-budget` (pages fully processed per similarity cluster, default 1).

Crawl a target and skip near-duplicate pages using TF-IDF similarity scoring to reduce noise on large sites.

```
$ katana -u https://example.com -pcs -pcsm tfidf -pcst 0.90
```

— Full flag set with defaults and a runnable example. v1.7.0

02 Secrets extractor in knowledge base NEW

85

The `-kb-secrets` flag enables a secrets extractor within the knowledge base that automatically detects exposed credentials and tokens during crawls. The `-kb-validate-secrets` flag validates detected secrets live against their respective providers by sending real API calls.

Crawl a target and extract secrets from the knowledge base, then validate them live against their providers.

```
$ katana -u https://example.com -kb -kb-secrets -kb-validate-secrets
```

— Named flags, mechanism, and runnable example given. v1.7.0

03 **Context field for lifecycle control** NEW 70

Adds a `Context` field to the `Options` struct for external lifecycle control — cancelling the provided `context.Context` stops the rate limiter, crawl session, queue, known-files requests, and headless browser.

— Named struct field with clear behaviour, no code sample. v1.7.0

04 **Endpoints extractor in knowledge base** NEW 60

The `-kb-endpoints` flag enables an endpoints extractor that classifies REST, GraphQL, SOAP, and XHR requests found during crawling.

— Named flag and classification types, but no example. v1.7.0

THINNER COVERAGE BELOW

05 **Headless engine callbacks** NEW 30

Adds callbacks to the headless engine, enabling programmatic interception of headless browser events.

— Describes capability but no API names or usage shown. v1.7.0

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

2 RELEASES · 2026-08-03 → 2026-08-04 BLOG ↗

GitGuardian Detection Engine scans code repositories to identify and alert on exposed secrets, credentials, and sensitive data in real-time.

> GitGuardian's Detection Engine expanded its secret-detection coverage with eight new detectors (Lovable, Slack Workflow Builder, n8n, Atlassian, AWS Cognito, OpenAI, Vercel OIDC, OpenCode) and a new Supabase JWT checker, while also tightening validation on GitHub, npm and PostgreSQL detectors and removing/disabling two detectors.

└──▷ WHAT SHIPPED · 10 FEATURES most completely described first ? what's the number?

01 **PostgreSQL credentials detector expanded matching** IMPROVED

70

PostgreSQL credentials detectors now match usernames containing dashes (e.g. `workflow-prod-user`) and raise the username length bound from 20 to 30 characters.

— Gives concrete example and exact numeric limit change

Detection Engine Updates Version 2.169 …

02 **New detectors for Slack, n8n, Atlassian and others** NEW

65

Seven new detectors were added: Slack Workflow Builder webhook trigger URLs, n8n API keys, Atlassian auth keys (OAuth 2.0 access tokens issued by `auth.atlassian.com`), AWS Cognito User Pool access tokens, OpenAI platform access tokens, Vercel OIDC workload-identity tokens, and OpenCode API keys.

— Names each detector type but no config/API to act on

Detection Engine Updates Version 2.169 …

03 **Azure Event Hub/Service Bus checker false-valid fix** IMPROVED

65

Azure Event Hub and Azure Service Bus checkers now only report a key as valid when the 200 response is a genuine Service Bus Atom feed, blocking false-valid results behind intercepting proxies.

— Explains exact validation mechanism and threat scenario fixed

Detection Engine Updates Version 2.169 …

THINNER COVERAGE BELOW

04 **GitHub token detector format validation** IMPROVED 50

The GitHub Personal Access Token, GitHub OAuth Token, and GitHub App Token detectors now skip values that fail format validation, reducing false positives from generic detectors.

— Names three specific detectors and the fix's effect

Detection Engine Updates Version 2.169.…

05 **Grafbase access token detector removed** BREAKING 50

The `grafbase_access_token` detector and checker have been removed entirely.

— Names exact detector removed, no migration guidance

Detection Engine Updates Version 2.169 …

06 **npm Token detector format validation** IMPROVED 45

The npm Token detector now skips values that fail format validation, reducing false positives from generic detectors.

— Clear mechanism but single narrow detector Detection Engine Updates Version 2.169.…

07 **Amadeus OAuth credentials detector disabled** BREAKING 45

The Amadeus OAuth credentials detector and checker have been disabled.

— Names detector disabled, no reason or migration path given

Detection Engine Updates Version 2.169 …

08 **Supabase service_role JWT checker** NEW 40

A new checker validates Supabase `service_role` JWTs.

— Names exact JWT type but no detail on validation logic

Detection Engine Updates Version 2.169 …

09 **Lovable token detectors** NEW 35

New detectors identify Lovable project access tokens and Lovable git tokens.

— Names the token types but no mechanism given Detection Engine Updates Version 2.169.…

10 **Generic Database Assignment detector port variant** IMPROVED 35

The Generic Database Assignment detector now incorporates the attached-port variant.

— Names the detector but mechanism is vague Detection Engine Updates Version 2.169 …

└──▷ BREAKING ON UPGRADE

! The `grafbase_access_token` detector and checker have been removed entirely.

! The Amadeus OAuth credentials detector and checker have been disabled.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sophos DNS Protection

Feedback ⓘ 1 RELEASE · 2026-08-05

NOTES ↗

Sophos DNS Protection is a security service that blocks malicious websites and threats by filtering DNS queries at the network level.

> Sophos DNS Protection adds a dedicated Generative AI web category, giving filtering policies granular control over AI services separate from other business categories.

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

01 **Generative AI filtering category** NEW 57

A new Generative AI category is added to the Business and generative AI category group (formerly Business-relevant site categories), letting filtering policies block or allow generative AI websites independently of other business categories.

— Names the category and group but no policy config path given

2026.32Released on August 5, 2026

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

Fastly is a content delivery network and edge computing platform that accelerates website performance and provides security services globally.

> Fastly's Next-Gen WAF now surfaces who created a given rule or signal, adding a `created_by` field to its API responses.

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

01 **created_by field in Next-Gen WAF APIs** NEW 60

The Next-Gen WAF Rules API and Next-Gen WAF Signals API responses now include a `created_by` field, identifying the email address of the user who created the rule or signal.

— Names exact field and two APIs but no example request/response shown.

Next-Gen WAF: `created_by` field now in…

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Docker Sandboxes

Feedback ⓘ 1 RELEASE · 2026-08-06 NOTES ↗

Docker Sandboxes provides isolated, lightweight environments for securely running and testing AI models and applications with controlled dependencies.

> Docker Sandboxes' v0.38.0 release adds local model routing for Claude Code, first-class MCP server management, enterprise proxy/NTLM/Kerberos support, per-sandbox network deny rules, and experimental GPU passthrough, while accompanying docs detail a new AI Governance Audit Logs feature (0.35.0+) and OpenCode Zen API key integration.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 OpenCode Zen API key custom secret with host proxy substitution

NEW 90

`sbx secret set-custom -g --host opencode.ai --env OPENCODE_API_KEY --value` stores OpenCode Zen API keys so the real key stays in the host secret store while the sandbox receives a placeholder that the host-side proxy substitutes on requests to `opencode.ai`. Sandboxes also need `sbx policy allow network opencode.ai:443` to reach the OpenCode Zen service.

Store an OpenCode Zen API key as a global custom secret so all OpenCode sandboxes can authenticate to opencode.ai without exposing the real key inside the sandbox.

```
$ sbx secret set-custom -g \
  --host opencode.ai \
  --env OPENCODE_API_KEY \
  --value "$OPENCODE_API_KEY"
```

Allow an OpenCode sandbox to reach the OpenCode Zen API over HTTPS after adding the custom secret.

```
$ sbx policy allow network opencode.ai:443
```

— Exact runnable commands and mechanism given product docs

02 AI Governance Audit Logs with local and cloud delivery NEW

85

Requires Docker Sandboxes `0.35.0`+ and an AI Governance subscription. The sandbox daemon writes local audit records as `.jsonl` files (JSON Lines) to per-OS paths —

`~/Library/Logs/com.docker.sandboxes/sandboxes/auditkit/` (macOS), `${XDG_STATE_HOME:-~/.local/state}/sandboxes/sandboxes/auditkit/` (Linux), `%LOCALAPPDATA%\DockerSandboxes\sandboxes\logs\auditkit\` (Windows) — named `audit-<utc-timestamp>-<process-uuid>-<seq>.jsonl`, finalized by atomic rename at rotation thresholds of 5 minutes, 1000 events, or 50 MiB. Records include evaluation events (resource, action, verdict, `deny_reason`) and session lifecycle events linked by `audit_session_id`, with fields like `audit_event_id`, `schema_version`, `username`, `org_id`, `hostname`, `network_egress`. Local delivery supports SIEM ingestion via log shippers (Splunk Universal Forwarder, Filebeat, CrowdStrike Falcon LogScale); Docker Cloud delivery (on by default when AI Governance is enabled) additionally powers a hosted audit log view, CSV export, and SIEM streaming from `app.docker.com`, configurable by org owners or custom-role users. Run `sbx policy ls` to confirm a `Governance: Managed by <org>` line, a prerequisite for audit records to be written.

Confirm that governance is active and audit logging is enabled before relying on local audit records.

```
$ sbx policy ls
```

— Full field/path/rotation detail plus one verification command, no enable steps

product docs

---

03 **Local model routing for Claude Code** NEW 85

`sbx run --model <name> claude` serves a local GGUF model via llmman, or an already-running Ollama install using an `ollama/` prefix, routing Claude Code away from `api.anthropic.com`. Requires enabling `sbx settings set platform.allowExperimentalFeatures true` and `sbx settings set feature.model true` first.

Route Claude Code at a local GGUF model instead of `api.anthropic.com` — useful for air-gapped or cost-sensitive environments.

```
$ sbx settings set platform.allowExperimentalFeatures true
sbx settings set feature.model true
sbx run --model llama3.2 claude
```

— Exact commands and enabling flags provided v0.38.0

---

04 **Enterprise proxy support with NTLM/Kerberos and TLS options**

NEW 85

Settings-driven upstream-proxy configuration via `proxy`, `proxy.sandbox`, `proxy.daemon`, and matching `no_proxy` keys defaults to the host OS system proxy

and proxies daemon traffic (image pulls, telemetry) too. `proxy.integratedAuth` enables NTLM/Kerberos/Negotiate proxy authentication on Windows, and `tls.allowNegativeSerial` supports TLS-inspecting proxies issuing certificates with negative serial numbers. Blocked HTTPS proxy connections now appear in `sbx policy log` even when the client aborts the TLS handshake.

Configure a corporate HTTPS proxy for all sandbox and daemon traffic with NTLM authentication on Windows.

```
$ sbx settings set proxy http://proxy.corp.example.com:8080
sbx settings set proxy.integratedAuth true
```

— Named settings keys with a runnable configuration example v0.38.0

05 **Per-sandbox network deny rules** NEW 65

`--deny-network HOST` flag on `sbx run` and `sbx create` records per-sandbox egress deny rules at creation time, e.g. to block a sandbox from reaching a specific host to prevent exfiltration.

Block a sandbox from reaching a specific host at creation time, e.g. to prevent exfiltration to an external service.

```
$ sbx run --deny-network api.example.com claude
```

— Exact flag with runnable example command v0.38.0

06 **First-class MCP server management with live index refresh**

IMPROVED 60

`sbx mcp` becomes a first-class, no-longer-hidden subcommand for MCP server management, with `sbx mcp add` supporting `--dir` and OAuth override flags; registering a server immediately refreshes the `mcp-find` / `mcp-add` index in every running sandbox without recreating them. Dynamic MCP tools `mcp-find`, `mcp-add`, and `mcp-config-set` are now also available in local sandbox mode.

— Named subcommand and flags but no worked example v0.38.0

07 **Experimental NVIDIA GPU passthrough** NEW 60

`sbx run --gpu` (experimental) enables NVIDIA VFIO GPU passthrough on Linux; enable with `sbx settings set feature.sandbox-gpu true`.

— Named flag and enabling setting, no worked walkthrough v0.38.0

THINNER COVERAGE BELOW

08 **sbx mount host bind-mount at matching path** IMPROVED

55

`sbx mount <sandbox> <host>` now bind-mounts the host path at the same path inside the sandbox when no container target is given; `:ro` mounts are enforced read-only at the virtio-fs host edge.

— Behaviour and enforcement point described, no example v0.38.0

09 **Secrets default to global scope** BREAKING

50

Service and custom secrets are now global by default, with a `--sandbox` flag for sandbox-scoped secrets; legacy positional and `--global` forms are deprecated with warnings.

— Names flag and deprecation but no migration example v0.38.0

10 **sbx mcp bundle command removed** BREAKING

40

The `sbx mcp bundle` command has been removed; single-server management via `sbx mcp add` / `ls` / `rm` is now the only supported flow.

— States removal and replacement, no migration steps v0.38.0

11 **sbx daemon restart command** NEW

40

`sbx daemon restart` stops and restarts the sandboxd daemon in the background.

— Command named but no further mechanism described v0.38.0

12 **Claude Vertex multi-region endpoint support** IMPROVED

40

Claude Vertex sandboxes now support `us` and `eu` multi-region endpoints.

— Names regions but no configuration steps given v0.38.0

13 **Configurable cloned workspace volume size** NEW

35

`DOCKER_SANDBOXES_CLONED_WORKSPACE_SIZE` environment variable configures the size of the cloned workspace volume.

— Env var named, no default or usage example given v0.38.0

14 **Homebrew stapled .dmg installer on macOS** IMPROVED

35

Homebrew installs now use a stapled `.dmg` artifact instead of a `.tar.gz` archive, improving Gatekeeper compatibility on macOS.

— Names artifact change, nothing for the reader to run v0.38.0

15 **sbx inspect shows custom secrets** IMPROVED

30

`sbx inspect` now displays custom secrets configured for a sandbox.

— Single-line change with no further detail v0.38.0

└──▷ **BREAKING ON UPGRADE**

! The `sbx mcp bundle` command has been removed; single-server management via `sbx mcp add` / `ls` / `rm` is the only supported flow.

! Service and custom secrets default to global scope; legacy positional and `--global` forms are deprecated with warnings.

WAS THIS USEFUL? 👍 👎

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

> OpenShell v0.0.99 adds a system CA root mode for sandboxes, starts honoring OCI image working directories, and fixes a `TCP_NODELAY` latency issue.

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

01 **System CA root mode for sandboxes** NEW 50

Adds a system CA root mode to the build, enabling sandboxes to trust the host's system certificate authorities.

— Names the capability and effect but no flag or config surface given. v0.0.99

02 **TCP_NODELAY latency fix** IMPROVED 30

The release includes a latency fix related to `TCP_NODELAY`, though no further explanation of the change is given.

— Names the socket option but lacks context on scope or impact. v0.0.99

03 **Honors OCI image working directories** IMPROVED 25

OpenShell now honors the working directory specified in OCI images, per the release summary; no further mechanism is described.

— Only mentioned briefly in the summary, no detail on behavior. v0.0.99

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Sigstore cosign

Feedback ⓘ 1 RELEASE · 2026-08-06 NOTES ↗

Code signing and transparency for containers and binaries

> cosign v3.1.3 adds support for signing OCI images with an X.509 certificate chain and auto-detects digest algorithms for public keys.

└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

01 **OCI signing with X.509 certificate chain** NEW 45

cosign now supports OCI signing with an X.509 certificate chain, enabling bring-your-own PKI workflows for container image signing.

— Names capability but no flag or command shown. v3.1.3

02 **Auto-detected digest algorithm for public keys** IMPROVED 40

cosign auto-detects the default digest algorithm for public keys, removing the need to manually specify it.

— Describes behaviour change but no flag or command detail. v3.1.3

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 1 RELEASE · 2026-08-06 NOTES ↗ CODE ↗

Prowler is the Open Cloud Security Platform trusted by thousands to automate security and compliance in any cloud environment. With thousands of ready-to-use checks and compliance frameworks, delivers real-time, customizable monitoring and seamless integrations, making cloud security simple, scalable, cost-effective for organizations of any size.

> Prowler 5.38.0 expands M365 compliance coverage with 12 new checks, adds 20 AWS privilege-escalation Attack Paths queries plus outcome-category metadata, and introduces Compliance Watchlists and multi-domain SAML SSO in Prowler Cloud/Private Cloud.

└──▷ WHAT SHIPPED · 10 FEATURES most completely described first ? what's the number?

01 **New M365 checks for CIS Foundations Benchmark v7.0.0** NEW

85

Adds 12 M365 provider checks: `admincenter_shared_bookings_disabled` (control 1.3.9); `defender_priority_account_protection_enabled` and `defender_strict_preset_security_policy_enabled` (controls 2.4.1, 2.4.2); `exchange_owa_mailbox_policy_personal_accounts_disabled` and `exchange_organization_reject_direct_send_enabled` (controls 6.3.2, 6.5.5); `teams_external_access_trial_tenants_blocked` (control 8.2.4); and six device-registration checks — `entra_device_registration_join_restricted`, `entra_device_registration_max_devices_per_user_limited`, `entra_device_registration_global_admins_not_local_admins`, `entra_device_registration_registering_user_not_local_admin`, `entra_device_registration_laps_enabled`, and `entra_policy_default_user_cannot_read_bitlocker_keys` (controls 5.1.4.x), all mapped to CIS Microsoft 365 Foundations Benchmark v7.0.0.

— Named checks and exact CIS control mappings given 5.38.0

02 **20 new AWS privilege-escalation queries in Attack Paths** NEW

80

Attack Paths adds 20 AWS privilege-escalation detection queries sourced from pathfinding.cloud, covering `iam:PassRole` escalations across AWS Batch, Braket, Cognito Identity, ECS, EMR, EMR Serverless, GameLift, Glue, EC2 Image Builder, Kinesis Analytics, HealthOmics, EventBridge Scheduler, SSM, and Step Functions, plus CodeDeploy/Step Functions existing-resource abuse, permissions-boundary removal with role assumption, and IAM Identity Center permission-set policy injection.

— Detailed services and escalation paths, no direct usage step 5.38.0

03 **Outcome category on Attack Paths query metadata** IMPROVED 70

Attack Paths query metadata now exposes an outcome category (Code execution, Privilege escalation, Public exposure, or Resource inventory) on the queries endpoint, enabling terminal outcome nodes in the graph.

— Names the endpoint and category values but no path/method 5.38.0

04 **Expanded privilege-escalation combinations in IAM policy check** IMPROVED 65

`iam_policy_allows_privilege_escalation` check gains 22 additional escalation combinations from pathfinding.cloud across AWS Batch, Braket, CodeDeploy, Cognito Identity, ECS, EMR, EMR Serverless, GameLift, Glue, EC2 Image Builder, Kinesis Analytics, HealthOmics, and EventBridge Scheduler.

— Named check and services covered, but no run instructions 5.38.0

05 **Multi-domain SAML SSO configuration** IMPROVED 65

SAML SSO now supports one primary email domain plus up to 19 additional domains under a single Identity Provider configuration and a single stable ACS URL (Prowler Cloud/Private Cloud).

— Exact domain limit and ACS URL detail, no config path 5.38.0

THINNER COVERAGE BELOW

06 **Compliance Watchlist for pinned frameworks** NEW 55

Compliance Watchlist (Prowler Cloud/Private Cloud) lets teams pin frameworks from any compliance view, filter the Compliance section to pinned frameworks only, and see pinned framework scores on the Overview page.

— Describes behaviour and UI location, no config key 5.38.0

07 **Sign-in method tags in Users table** NEW 50

Users table in Prowler Cloud now shows each account's sign-in methods as tags, including email/password, Google, GitHub, SAML with linked domains, and Partner SSO.

— UI location given, no further mechanism 5.38.0

08 **SBOMs and build-provenance attestations for container images** NEW 50

Prowler API, UI, SDK, and MCP container images now publish per-architecture SBOMs and build-provenance attestations as OCI attestations.

— Names artifacts and format but no retrieval command 5.38.0

09 **Grype scanning alongside Trivy in image security gates** IMPROVED

40

Grype now runs alongside Trivy across container-image security gates, blocking on fixable high and critical findings.

— Names tools and gate behaviour, no config surface 5.38.0

10 **Checksum verification for downloaded build binaries** IMPROVED

35

API and SDK container builds now verify checksums of downloaded PowerShell, Trivy, and zizmor binaries before installation.

— Names binaries but process is internal, not user-actionable 5.38.0

WAS THIS USEFUL? 👍 👎

Wiz

Feedback ⓘ 1 RELEASE · 2026-08-06 BLOG ↗

Wiz is a cloud security platform that identifies and remediates risks across cloud infrastructure and applications.

> Wiz added automated DISA STIG compliance assessment for Amazon Linux 2023 and Windows Server 2025 in AWS GovCloud, closing gaps where official SCAP content isn't yet available.

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

01 **Automated DISA STIG assessment for AL2023 and Windows Server 2025**

NEW

68

Wiz adds automated DISA STIG compliance assessment covering Amazon Linux 2023 (`AL2023 V1R2` and `V1R3` benchmarks) and Windows Server 2025 , filling gaps where official SCAP content is not yet available. It delivers continuous hardening validation for federal agencies and defense contractors operating in AWS GovCloud, supporting FedRAMP Rev5, FedRAMP 20x, CMMC, and DoD authorization requirements.

— Names benchmarks and frameworks but no config or API surface given. 2026-08-06-c676879b

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Wiz Cloud

Feedback ⓘ 1 RELEASE · 2026-08-06 BLOG ↗

Wiz Cloud is a cloud security platform that identifies and remediates risks across cloud infrastructure, applications, and data.

> Wiz added automated DISA STIG compliance assessment for Amazon Linux 2023 and Windows Server 2025, surfacing agentless host-level findings directly in the Compliance Posture dashboard.

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

01 DISA STIG Amazon Linux 2023 assessment via HCR NEW

80

Adds DISA STIG Amazon Linux 2023 benchmark support via Host Configuration Rules (HCR), providing continuous agentless assessment across all connected cloud instances without SCAP content or scan credentials. Enable it in the Wiz portal under `Compliance Posture > Manage Frameworks` by searching for `DISA STIG Amazon Linux 2023` — assessment begins automatically for all connected instances.

Enable the DISA STIG Amazon Linux 2023 framework to start continuous agentless assessment across your entire AL2023 fleet.

> 📍 In the Wiz portal, go to Compliance Posture › Manage Frameworks, search for `DISA STIG Amazon Linux 2023`, and enable the benchmark — assessment begins automatically for all connected instances.

— Names HCR mechanism and exact portal path, no API. 2026-08-06-d1ea617b

02 DISA STIG Windows Server 2025 assessment via HCR NEW

77

Adds DISA STIG Microsoft Windows Server 2025 benchmark support via Host Configuration Rules (HCR), with an initial automated rule set and expanding coverage planned in upcoming releases. Enable it via `Compliance Posture > Manage Frameworks` by searching for `DISA STIG Microsoft Windows Server 2025` — new instances are assessed automatically on deployment.

Enable the DISA STIG Microsoft Windows Server 2025 framework to continuously assess Windows Server 2025 hosts as they are deployed.

> 📍 In the Wiz portal, go to Compliance Posture › Manage Frameworks, search for `DISA STIG Microsoft Windows Server 2025`, and enable the benchmark — new instances are assessed automatically on deployment.

— Exact portal steps given; rule set breadth unspecified. 2026-08-06-d1ea617b

THINNER COVERAGE BELOW

03 **STIG findings in Compliance Posture dashboard** NEW 50

STIG findings for both new frameworks surface in the Compliance Posture dashboard alongside existing frameworks, with per-rule pass/fail status per instance and rollup into category and subcategory scores.

— Describes rollup behaviour but no exact dashboard path. 2026-08-06-d1ea617b

04 **Filter findings by STIG severity category** NEW 45

Supports filtering findings by STIG severity category (`CAT I` , `CAT II` , `CAT III`) to prioritize high-impact failures.

— Names the three categories, no filter UI detail. 2026-08-06-d1ea617b

05 **Compliance trend, reports and automation for STIG checks** NEW 45

Enables Compliance Posture Trend tracking, auditor report generation, and automation rules to create tickets or trigger remediation workflows for failed STIG checks.

— Lists linked capabilities without configuration specifics. 2026-08-06-d1ea617b

06 **STIG assessment available under existing licenses** NEW 25

Available to all Wiz for Government customers under existing Wiz licenses with no additional license required.

— Simple licensing note with no mechanism or steps. 2026-08-06-d1ea617b

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 3 RELEASES · SEEN 2026-08-06 NOTES ↗

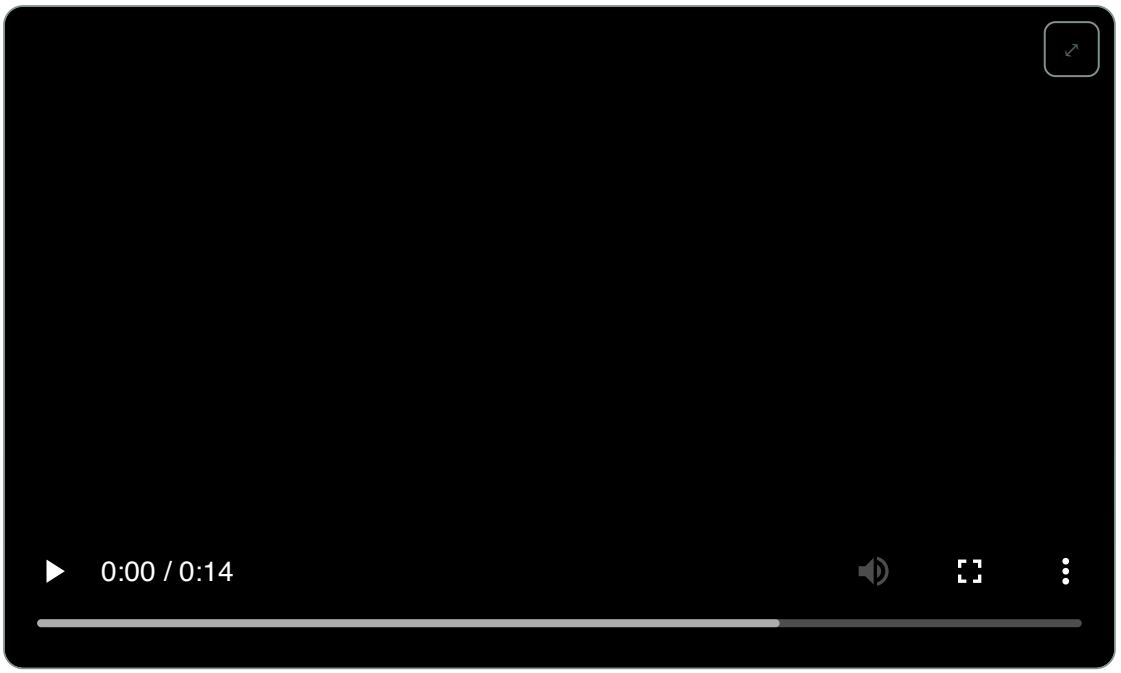
> Iru shipped three Apple device management updates: Apple beta program enrollment via the Software Update Library Item, phased rollout support extended to Managed OS updates, and Authenticated Guest Mode for Platform SSO on macOS.

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

01 **Apple beta program enrollment via Software Update Library Item**

NEW 60

The Software Update Library Item now supports Apple beta program enrollment, allowing end users to opt-in or letting admins enforce a specific beta program without tying it to a specific version.



— Names the library item and behavior but no config keys.

Offer Apple beta releases with the Soft…

THINNER COVERAGE BELOW

02 **Phased rollouts for Managed OS updates** NEW 50

Extends phased rollout support — previously available for Mac Auto Apps — to Apple Managed OS, letting admins stage OS update declarations across their device fleet instead of pushing to all devices simultaneously.

— Explains mechanism and scope but no config surface named.

Phased Rollouts for Managed OS for Appl…

03 Authenticated Guest Mode for Platform SSO on Mac NEW

40

Adds Authenticated Guest Mode configuration support to the Single Sign-On Extension Library Item when using Platform SSO with macOS.

— Names the library item but lacks detail on how it works.

Platform SSO for Mac: Authenticated gue…

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ 2 RELEASES · 2026-08-06 BLOG ↗

Tenzir is a security data platform that ingests, transforms, and analyzes security telemetry and threat intelligence data.

> Tenzir's v6.9.0 release line adds native Apache Iceberg output, automatic NetFlow/IPFIX decoding, and OCSF 1.9.0 schema support, alongside several new configuration keys for faster, lower-memory catalog startup and data-integrity validation.

└──▷ WHAT SHIPPED · 7 FEATURES most completely described first ? what's the number?

01 Native Apache Iceberg output via to_iceberg NEW

95

A new `to_iceberg` operator writes events directly into Apache Iceberg tables through a REST catalog, with `mode` (`create_append`, `create`, `append`), `partition_by` supporting hidden-partition transforms (`year`, `month`, `day`, `hour`, `bucket`, `truncate`), plus `max_size`, `timeout`, and `buffer_size` for file rotation. It writes zstd-compressed Parquet with automatic schema evolution for heterogeneous streams, and connects to S3-compatible stores via `aws_iam`, `s3_endpoint`, `s3_path_style`, bearer-token catalogs via `token`, AWS Glue/Amazon S3 Tables via `catalog_aws_service`, and Google BigLake via `gcp_service_account_key` or `gcp_auth`.

Write a live OCSF stream to an Iceberg table partitioned by class and day, with automatic schema evolution.

```
$ subscribe "ocsf"
ocsf::cast
to_iceberg "security.ocsf", catalog =
"https://catalog.example.com", partition_by = [class_uid,
day(time)]
```

Write a live OCSF stream to an Apache Iceberg table partitioned by class and day, creating the table automatically if it doesn't exist.


```
$ subscribe "ocsf"
ocsf::cast
to_iceberg "security.ocsf", catalog =
"https://catalog.example.com", partition_by = [class_uid,
day(time)]
```

— Full option set and a runnable pipeline example given.

Tenzir Node v6.9.0-rc.2: Native Apache … | Tenzir Node v6.9.0: Native Apache Icebe…

02 **Faster catalog startup via index tuning config keys** NEW

95

New config keys reduce startup time and memory on nodes with many partitions: `tenzir.index.load-concurrency` tunes worker threads for concurrent partition synopsis loading (defaults to hardware concurrency); `tenzir.index.lazy-sketches` defers Bloom-filter sketch loading to query time, bounding warm memory via an LRU cache sized by `tenzir.index.sketch-cache-bytes` (default 1 GiB); and `tenzir.index.skip-synopsis-verification` skips recursive FlatBuffers verification of partition synopses at startup, trading corruption robustness for speed on trusted networked storage such as NFS.

Reduce startup memory and load time on a node with many partitions backed by NFS.

yaml

```yaml
tenzir.index.load-concurrency: 16
tenzir.index.lazy-sketches: true
tenzir.index.sketch-cache-bytes: 2147483648
tenzir.index.skip-synopsis-verification: true
```

Reduce startup memory and time on a node with many partitions by deferring Bloom-filter sketch loading and skipping synopsis verification on trusted NFS storage.

yaml

```yaml
tenzir.index.lazy-sketches: true
tenzir.index.skip-synopsis-verification: true
tenzir.index.sketch-cache-bytes: 2147483648  # 2 GiB
tenzir.index.load-concurrency: 16
```

— Named config keys with defaults and a runnable config example.

Tenzir Node v6.9.0-rc.2: Native Apache … | Tenzir Node v6.9.0: Native Apache Icebe…

03 **Automatic NetFlow and IPFIX decoding** NEW

90

A new `read_netflow` operator automatically decodes NetFlow v5, NetFlow v9, and IPFIX from byte streams and UDP envelopes in a single pipeline, learning templates per exporter and preserving unknown elements as blobs, with no manual parser configuration or protocol pre-selection needed.

Ingest mixed NetFlow v5/v9 and IPFIX from UDP on port 2055 in a single pipeline.

```
$ accept_udp "0.0.0.0:2055", binary = true
read_netflow
```

Ingest mixed NetFlow v5/v9 and IPFIX traffic from a UDP listener without pre-selecting a protocol version.


```
$ accept_udp "0.0.0.0:2055", binary = true
read_netflow
```


— Names operator, versions handled, and a runnable UDP example.

Tenzir Node v6.9.0-rc.2: Native Apache … | Tenzir Node v6.9.0: Native Apache Icebe…

04 **Store batch validation for corruption diagnosis** NEW 60

The `tenzir.validate-store-batches` config key fully validates every Arrow record batch read from a persisted Feather store, helping diagnose on-disk corruption including malformed UTF-8; disabled by default.

— Named config key and default, no usage example.

Tenzir Node v6.9.0-rc.2: Native Apache … | Tenzir Node v6.9.0: Native Apache Icebe…

05 **Faster ocsf::cast variant encoding and null-filling** IMPROVED 60

`ocsf::cast` with `encode_variants=true` now encodes free-form objects such as `unmapped` significantly faster, lowering CPU usage for high-volume streams, and `ocsf::cast` with `null_fill=true` now uses significantly less memory and CPU by avoiding per-batch column rebuilds.

— Named operator and flags but no example command.

Tenzir Node v6.9.0-rc.2: Native Apache …

THINNER COVERAGE BELOW

06 **Dynamic field assignment with expression keys** NEW 45

Fields can now be assigned using expression keys via `result[key] = value`, letting field names be determined per-event at runtime instead of being fixed at pipeline-definition time.

— Syntax named but no runnable example given. Tenzir Node v6.9.0: Native Apache Icebe…

07 **OCSF 1.9.0 schema support** NEW 40

Adds support for OCSF 1.9.0 (and all earlier bundled stable versions), plus preview support for OCSF `1.10.0-dev`.

– Version numbers named but no usage detail. Tenzir Node v6.9.0-rc.2: Native Apache … Tenzir Node v6.9.0: Native Apache Icebe…

WAS THIS USEFUL? 👍 👎

Beelzebub

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

A secure low code deception runtime framework, leveraging AI for System Virtualization.

> Beelzebub v3.9.0 shifts plugin management to declarative YAML configuration, adds boot-time background plugins, and introduces shared JSON Schema validation for honeypot configs.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 **ServicePlugin background plugins at boot** NEW 50

Adds `ServicePlugin` support for background plugins that start automatically at boot time.

– Names concrete plugin type but no usage steps v3.9.0

02 **Shared JSON Schema validation for honeypot configs** NEW 43

Adds shared JSON Schema validation for honeypot configurations, catching config errors before deployment.

– Names format and benefit, lacks schema location or command v3.9.0

03 **YAML-based plugin declarations** IMPROVED 33

Plugin declarations move to YAML config files, enabling declarative plugin management via configuration files instead of the prior approach.

– Names config format but no key or file path given v3.9.0

04 **Improved installer experience** IMPROVED 5

Improves the installer experience.

– No detail on what changed or how v3.9.0

WAS THIS USEFUL? 👍 👎

◆ SIEM

Sumo Logic

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

Sumo Logic is a cloud-based monitoring and analytics platform for logs, metrics, and security data.

> Sumo Logic added Mobot, a chat-based assistant that builds and edits Automation Service playbooks from natural-language descriptions.

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

01 **Mobot natural-language playbook creation** NEW 60

Mobot supports playbook creation and editing via natural-language chat: you describe the automation you want, Mobot checks available integrations, asks clarifying questions, proposes a plan, configures each step, and saves the result as a draft for review and publishing.

— Describes full workflow but no UI path or config named

Manage Playbooks with Mobot (Automation…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

1 RELEASE · 2026-08-07 BLOG ↗

Feedly Threat Intelligence aggregates security feeds and threat data to help organizations monitor emerging cybersecurity threats and vulnerabilities.

> Feedly's Vulnerability Intel Agent added a dedicated Atlassian security release filter with product, attack-path, and patch-policy filtering, alerting, and expanded vendor advisory coverage.

└──▷ WHAT SHIPPED · 7 FEATURES most completely described first ? what's the number?

01 **Attack path filtering for CVEs** NEW 65

CVEs can be filtered by attack path — attack vector, privileges required, and user interaction — to isolate unauthenticated, no-interaction reachable vulnerabilities.

— Names specific attack-path dimensions used for isolation.

Enrich and triage Atlassian security re…

02 **Patch policy threshold filtering (CVSS/EPSS)** NEW 65

Filtering by patch policy thresholds supports CVSS score cutoffs, EPSS probability cutoffs, or both combined, letting teams isolate CVEs that breach both thresholds simultaneously to match internal patch policies.

— Names exact scoring metrics and combined-threshold behavior.

Enrich and triage Atlassian security re…

03 **Latest Atlassian Security Release filter** NEW 60

A Latest Atlassian Security Release filter in the Vulnerability Intel Agent always points to the most recent Atlassian advisory, ingested within minutes of publication, eliminating the need for manual updates each cycle.

— Names the exact filter but not full UI steps. Enrich and triage Atlassian security re…

04 **Product-level CVE filtering for Atlassian stack** NEW 60

Filtering supports scoping Atlassian CVEs down to specific products in a team's stack, such as Confluence Data Center, Sourcetree, and Jira, to surface only relevant exposures.

— Names concrete products but not exact filter mechanics.

Enrich and triage Atlassian security re…

THINNER COVERAGE BELOW

05 **Expanded vendor advisory coverage** IMPROVED 45

Vendor advisory support in the Vulnerability Intel Agent now includes Adobe, Android, Apache, Apple, Microsoft, Oracle, and Atlassian.

— Lists named vendors but no usage mechanism. Enrich and triage Atlassian security re…

06 **Alerting on filtered CVE matches** NEW 35

Alerts are available via a deployed Intel Agent, notifying teams when newly published CVEs match their configured filters.

— Mechanism described briefly, no setup detail. Enrich and triage Atlassian security re…

07 **Configurable enrichment columns in output table** IMPROVED 35

Enrichment columns in the Vulnerability Intel Agent's output table are reorderable and individually showable or hidable to match team triage workflows.

— Describes UI capability without exact navigation path.

Enrich and triage Atlassian security re…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Tenable

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

Tenable is a cybersecurity platform that identifies, measures, and remediates vulnerabilities across an organization's entire attack surface.

> Tenable introduced Hexa AI, an agentic capability within Tenable One that automates exposure remediation end-to-end using intent-driven Routines and human-approval guardrails.

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

01 **Intent-driven Routines for remediation workflow** NEW

55

Routines encode the multi-step remediation workflow — asset scoping, policy targeting, patch execution, status check, and re-scan confirmation — into a single automated sequence triggered by intent.

— Lists workflow steps but no exact command or config to invoke 2026-08-05-bfc66599

02 **Tenable Hexa AI agentic remediation** NEW

48

Tenable Hexa AI is an agentic capability within Tenable One that automates the full remediation cycle — scoping, patch deployment, and verification — across integrated platforms such as Jamf.

— Describes scope and integration but no config or API surface 2026-08-05-bfc66599

03 **Human approval guardrails on AI autonomy** NEW

35

AI autonomy is bound strictly to defined user permissions within the Tenable One harness, with human approval gates required before changes are applied.

— Explains safety mechanism conceptually, no configurable permission details given 2026-08-05-bfc66599

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Tenable One

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

Tenable One is an integrated cybersecurity platform that unifies vulnerability management, cloud security, and compliance across an organization's entire attack surface.

> Tenable One introduces Tenable Hexa AI, an agentic engine that turns natural-language objectives into gated, end-to-end vulnerability remediation routines.

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

01 **Natural-language intent-driven routines with guardrails** NEW 70

Intent-driven routines let practitioners define an objective, guardrails (e.g., 'Never launch a credentialed scan against anything tagged OT,' 'Open no more than twenty-five tickets in a run'), and a cadence (on-demand or scheduled) in plain natural language — no scripts required.

— Concrete guardrail examples give a clear starting point 2026-08-05-1fa9e5be

02 **Proposal-gated execution before changes** NEW 60

Before writing any changes, Hexa presents a plan showing device count, device list, policy, deployment window, and blast radius, and requires explicit approval before proceeding.

— Detailed plan fields but no exact command to trigger it 2026-08-05-1fa9e5be

THINNER COVERAGE BELOW

03 **Hexa AI agentic remediation engine** NEW 55

The new Tenable Hexa AI agentic engine automates end-to-end vulnerability remediation within a single conversational workflow — enumerating affected assets, mapping CVEs to fix versions, and executing patch deployments without manual tool-switching.

— Explains mechanism but no concrete UI path or command 2026-08-05-1fa9e5be

04 **Mid-workflow scope narrowing** NEW 40

Supports scope narrowing mid-workflow — tell Hexa to restrict enumeration to a single business unit and it re-scopes and returns a revised plan before executing.

— Describes behavior with a concrete example but no exact syntax 2026-08-05-1fa9e5be

05 **Routines adapt to infrastructure changes** IMPROVED 35

Routines adapt automatically to infrastructure changes such as new asset classes, renamed tags, or failed scans, without requiring script updates.

— States behavior but no mechanism detail 2026-08-05-1fa9e5be

06 **Explicit conflict surfacing between objective and guardrails**

NEW

25

Hexa surfaces conflicts between an objective and its guardrails explicitly rather than silently dropping either constraint.

— Brief prose description, no mechanism or UI detail 2026-08-05-1fa9e5be

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

iDenfy

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06 NOTES ↗

iDenfy is an identity verification platform that performs KYC/AML checks through document scanning and biometric analysis for compliance.

> iDenfy's headline shipment this window is Age Estimation — a selfie-only age-gating capability with automatic document step-up, webhook delivery, and a full dashboard for session configuration and review — alongside new Bank Verification, Face Authentication, Fraud Prevention, AML Screening and KYB APIs, plus a long tail of KYB/IDV language, workflow, billing, and SDK improvements.

└──▷ WHAT SHIPPED · 44 FEATURES most completely described first ? what's the number?

01 **Age Estimation session management dashboard** NEW 96

New Age Estimation section in the dashboard for browsing, filtering, searching and inspecting sessions, with a New estimation form exposing Minimum age (1–120), Safety margin (0–20 years), ID verification for uncertain results, Retry limit (1–3), Save selfie photo, and Link expiry time across General settings and Verification redirects tabs. Supports a per-session Webhook URL override and per-outcome redirect URLs (Success, Underage, Uncertain); the session list groups results into Results, Needs Attention, In Progress, Expired with filtering by client ID/creation date and search by session ID or client reference; session detail view surfaces a Step-up link to the underlying identity verification. Step-up Document Scans bill from Age Estimation finances rather than counting against document scan usage; viewing requires the View age estimation data permission and creating sessions requires `MANAGE_AGE_ESTIMATION`.

— Exact config keys, ranges, permission names and nav given snapshot-20260806

02 **Age Estimation capability, API and webhooks** NEW 83

Selfie-only minimum-age check that avoids a document scan in the common case; triggers an automatic document step-up to read an exact date of birth from the ID and compare the ID portrait against the capture selfie when the AI model lacks confidence. Session outcomes `SUCCESS`, `UNDERAGE`, `UNCERTAIN`, `FACE_MISMATCH`, `ATTEMPTS_EXCEEDED` are delivered via a signed Age Estimation Webhooks HTTP POST with idempotent delivery as soon as a session finishes; lifecycle statuses `PENDING`, `PROCESSING`, `COMPLETED`, `FAILED`, `EXPIRED` are returned by the Partner API, which pre-checks finances and rejects requests with `402` when funds are insufficient. Billed from a dedicated Age Estimation finances pool separate from standard KYC charges.

— Names outcomes, statuses, 402 behavior; lacks endpoint paths snapshot-20260806

03 **Branded email templates for KYB workflows** NEW 77

New Email template configuration under Business Verifications → Configurations → Email template lets teams create branded emails for three template types — a new business verification session URL, a renewed session URL requesting an information update, and a completed-company verification result — with an Add verification link action that inserts a session URL placeholder into the email body.

— Exact UI path and template types given, no API surface product docs

04 **KYB API for company verification** NEW 62

New KYB API enables company verification via registry data retrieval, beneficial owner screening, and AML checks, with dashboard-based compliance reviews. A Stakeholder Roles reference documents five roles — CEO, Representative, Shareholder, UBO, and ABO — including mappings for partnerships, foundations, and holding chains.

— Names roles and checks, no endpoint paths given product docs

THINNER COVERAGE BELOW

05 **Fraud Prevention API** NEW 50

New Fraud Prevention API adds AI risk scoring, VPN and proxy detection, phone validation, and address verification checks.

— Lists check types, no endpoint or scoring detail product docs

06 **AML Screening API** NEW 50

New AML Screening API supports sanctions, PEP, and adverse media checks, offered as single checks or ongoing monitoring.

— Lists check categories, no endpoint or config product docs

07 **Android SDK screenshot blocking, camera detection, and 3D liveness**

IMPROVED 50

Android SDK 8.7.1 adds screenshot blocking and virtual camera detection; Android SDK 8.7.6 adds optional 3D liveness and reduces SDK size.

— Names exact SDK versions and features snapshot-20260806

08 **Bank Verification API** NEW 47

New Bank Verification API verifies customer bank accounts and financial information programmatically, with session creation and webhook callbacks.

— Describes flow and webhook, no endpoint path product docs

09 **Face Authentication API** NEW 47

New Face Authentication API supports biometric face matching to re-authenticate returning users during account login, high-risk actions, and periodic checks.

— Names use cases, no endpoint or session model `product docs`

10 **Tax identification fields in KYB** NEW 41

Dedicated tax identification fields — VAT Number, `EIN`, `TIN` — added to KYB workflows.

— Names exact fields, no form location given `snapshot-20260806`

11 **Bank Card Verification step in KYC flow** NEW 38

New Bank Card Verification step confirms that the card used at onboarding belongs to the verified user, added as an additional step in the KYC flow.

— Describes purpose, no config or trigger detail `snapshot-20260806`

12 **Supported Languages reference (36 languages)** NEW 38

New Supported Languages reference documents 36 languages, including Latin and non-Latin scripts, with full locale codes for KYC and KYB flows.

— Gives a concrete count, but no lookup mechanism `product docs`

13 **Risk assessment in company verification API** NEW 35

The company verification API now returns risk assessment results as part of KYB responses.

— No endpoint name or response field detail given `snapshot-20260806`

14 **External Reference field in KYC reports** NEW 35

External Reference field added to KYC verification reports.

— Names the field, no usage detail `snapshot-20260806`

15 **Blocklist API endpoint** NEW 34

New Blocklist API endpoint for managing blocklisted identities.

— Names the endpoint but not its path or method `snapshot-20260806`

16 **KYB automation re-run endpoint** NEW 34

New API endpoint to re-run KYB automation on existing records.

— Names capability, no path or method given `snapshot-20260806`

17 **Dedicated accounting email for PAYG partners** NEW 33

Pay-As-You-Go partners can set a dedicated accounting email address for invoices and payment receipts.

— Describes what it does, no UI path named `snapshot-20260806`

18 **IDV mobile bottom sheet layout option** NEW 32

Identity verification (IDV) flows on mobile gain a bottom sheet layout option.

— Names UI option, no toggle location given `product docs`

19 **Configurable KYB company expiration period** NEW 30

KYB company verification records now support a configurable expiration period.

— Names the setting, no default value or location `snapshot-20260806`

20 **KYC API documentation for document checks and review** IMPROVED 30

Identity Verification (KYC) API documentation now covers document checks, liveness detection, and AI plus human review.

— Documentation scope only, no new mechanism described `product docs`

21 **New Digital ID support: Austrian Handy Signatur, Freja eID, Mobile ID**

NEW 30

New Digital ID support added: Austrian Handy Signatur, Freja eID, Mobile ID (Sprint 126).

— Names three schemes, no integration detail `snapshot-20260806`

22 **Companies House integration** NEW 29

Companies House integration added (Sprint 125) for company data lookups.

— Names integration, no API or field detail `snapshot-20260806`

23 **Risk Assessment PDF export** NEW 28

Risk Assessment PDF export now available.

— Names the export, no trigger or location given `snapshot-20260806`

24 **Custom redirect URL for iFrame** NEW 28

Custom redirect URL for iFrame added (Sprint 125).

— Names config surface, no key or field name `snapshot-20260806`

25 **KYB shared identity verification across companies** NEW 27

KYB now supports shared identity verification across multiple companies.

– Bare description, no mechanism explained snapshot-20260806

26 **KYB Verify Later automated IDV reminders** NEW 27

KYB Verify Later feature introduces automated IDV reminders.

– Names feature, no cadence or config given snapshot-20260806

27 **iOS SDK NFC improvements** IMPROVED 27

iOS SDK 8.7.7 adds NFC improvements.

– Names version, no specifics on the improvement snapshot-20260806

28 **New digital IDs for India and Philippines** NEW 25

New digital ID support added for India and the Philippines.

– Names countries but no scheme names or mechanism snapshot-20260806

29 **Brazil Digital CNH document support** NEW 25

Brazil Digital CNH (driver's license) document support added.

– Names document type, no scope or flow detail snapshot-20260806

30 **Auto document type detection** NEW 23

Auto document type detection for direct processing added (Sprint 126).

– Bare description of the feature snapshot-20260806

31 **PAYG social login** NEW 21

Pay-As-You-Go partners can now offer social login during onboarding.

– Bare description, no providers or config named snapshot-20260806

32 **AI reviewer spelling variation matching** IMPROVED 21

The AI reviewer gains spelling variation matching.

– Bare description, no scope or examples snapshot-20260806

33 **Korean and Traditional Chinese IDV language support** NEW 21

Korean and Traditional Chinese language support added to IDV (Sprint 124).

– Names languages, no locale codes given snapshot-20260806

34 **SMS verification** NEW — 21

SMS verification added (Sprint 123).

— Bare feature name, no flow detail snapshot-20260806

35 **Arabic language and RTL support for IDV** NEW — 21

Arabic language and RTL support added to IDV (Sprint 127).

— Names language/RTL, no config detail snapshot-20260806

36 **KYC multi-theme branding support** NEW — 21

KYC multi-theme branding support added (Sprint 127).

— Named feature, no theming mechanism snapshot-20260806

37 **KYB sole proprietorship identity verification** NEW — 21

KYB sole proprietorship identity verification added (Sprint 127).

— Names entity type, no flow detail snapshot-20260806

38 **AI KYB Reviewer fraud prevention capabilities** IMPROVED — 21

AI KYB Reviewer gains fraud prevention capabilities (Sprint 122).

— No mechanism or scope described snapshot-20260806

39 **Spanish language support for KYB** NEW — 19

Business verification (KYB) now supports the Spanish language.

— Single sentence, no config or scope detail product docs

40 **PNG logo uploads** NEW — 19

PNG logo uploads now supported.

— One-line addition, no location named snapshot-20260806

41 **KYB Dynamic Workflows** NEW — 19

KYB Dynamic Workflows introduced (Sprint 124).

— Named but undescribed feature snapshot-20260806

42 **Dashboard records open in new tab** IMPROVED — 19

Dashboard records can now be opened in a new browser tab.

— Small UI tweak, minimally described snapshot-20260806

43 **Turkish translation improvements for IDV** IMPROVED — 18

Turkish-language translations for IDV flows were improved.

— No specifics on what changed `product docs`

44 **Watermarks on PDFs** NEW 15

Watermarks on PDFs added (Sprint 122).

— One-line addition, no configuration detail `snapshot-20260806`

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

Clerk is a user authentication and identity management platform that provides sign-up, login, and user profile management for web and mobile applications.

> Clerk added Device Trust to require a second factor on new-device password sign-ins, introduced Google Workspace Directory Sync with role mapping, and published community-maintained Flutter integration docs.

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

01 **Google Workspace Directory Sync** NEW 43

Syncs users and groups directly from Google Workspace into Clerk, with support for role mapping and custom attributes.

— Names role mapping and custom attributes but no setup steps

`Google Workspace Directory Sync (2026-0…`

02 **Device Trust for new-device sign-ins** NEW 38

A new Device Trust capability automatically requires a second factor when a user signs in with a password from a new device and has not enabled MFA.

— Explains behaviour but no config surface or flag named `product docs`

03 **Community-maintained Flutter integration docs** NEW 23

Adds community-maintained Flutter integration support, with dedicated documentation for integrating Clerk into Flutter applications.

— Only states docs exist, no specific API or steps given `product docs`

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ 1 RELEASE · 2026-07-22 BLOG ↗

SailPoint is an identity governance and access management platform that automates user access provisioning, certification, and compliance across enterprise systems.

> SailPoint introduced Agentic Fabric, a full AI agent and non-human identity governance layer for Identity Security Cloud spanning endpoint and browser sensors, SIEM/EDR and cloud source connections, a unified Non-Human Identity Registry, sanctioned/unsanctioned app classification, identity graphs, audit reporting, and configurable aggregation, alongside new Dark Mode support across Agentic Fabric and select Identity Security Cloud pages.

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

01 **Endpoint agent sensor for AI discovery** NEW 70

A background endpoint agent, managed by navigating to `Admin > Global > Agent Settings > Sensors` and selecting the `Endpoint Agent` tab, discovers AI agent software installed on managed laptops and desktops.

After activating Agentic Fabric, manage the endpoint agent sensor that discovers AI software running on corporate endpoints.

> 📍 In Identity Security Cloud, go to Admin › Global › Agent Settings › Sensors, then select the Endpoint Agent tab to view and update endpoint agent configuration.

— Exact admin path given; describes purpose but not internals. product docs

02 **SIEM and EDR connections for agent monitoring** NEW 68

Navigating to `Admin > Global > Agent Settings` and selecting `SIEM and EDR Connections` lets admins connect EDR or SIEM platforms so Agentic Fabric can discover and monitor AI agents.

Connect an EDR or SIEM platform post-activation so Agentic Fabric can automatically find available data sources for AI agent monitoring.

> 📍 In Identity Security Cloud, go to Admin › Global › Agent Settings, then select SIEM and EDR Connections from the left panel to add or update your EDR/SIEM integration.

— Exact nav path and example steps, but UI-only action. product docs

03 **Identity provider correlation for AI agent ownership** NEW 68

Under `Admin > Security Settings > Service Provider`, admins manage identity providers (Microsoft Entra ID or Okta) post-activation to correlate human identities to the AI agents they own.

— Named path and providers, but no example walkthrough. product docs

04 **Business Apps sanction status grouping** NEW 68

A new Business Apps grouping lets administrators declare a sanction status — sanctioned, unsanctioned, or unknown — for logical applications and agents, with matching agents automatically inheriting that classification.

Classify which AI tools and agents are permitted or prohibited so that matching agents automatically inherit the sanctioned/unsanctioned status.

> 📍 In Identity Security Cloud, go to Agentic Fabric › Business Apps, select a business app, and set its sanction status to Sanctioned, Unsanctioned, or Unknown.

— Exact statuses and example steps, though UI-only. product docs

05 **Cloud platform sources for non-human identity aggregation** NEW 65

Via `Admin > Connections > Sources`, admins connect cloud platforms (AWS, Azure, and other AI platforms) as sources for aggregating non-human identities.

— Named path and platforms, no example given. product docs

06 **Browser extension sensor for GenAI governance** NEW 60

A browser extension sensor, managed via `Admin > Global > Agent Settings > Sensors` under the `Browser Extension` tab, governs GenAI usage across managed and unmanaged SaaS environments.

— Named admin path but no usage example provided. product docs

THINNER COVERAGE BELOW

07 **Configurable per-dataset aggregation for non-human identities** NEW 57

A datasets and resources model lets SailPoint sources aggregate non-human identities, AI agents, MCP servers, credentials, and IAM role objects on a configurable per-dataset aggregation schedule.

— Names object types and schedule concept, no exact config key. product docs

08 **Non-Human Identity Trends widget** NEW 55

A new widget on the `MySailPoint` home page shows total non-human identity counts over a selected time window, with day-level detail available on hover.

— Describes behaviour but no interaction steps. product docs

09 **Non-Human Identity Registry** NEW 50

A new unified Non-Human Identity Registry centralizes discovery and management of agents, applications, accounts, MCP servers, credentials, and endpoints.

— Names object types covered but no navigation or config given. product docs

10 **Audit and compliance reporting for AI agents** NEW 45

Audit and compliance reporting produces auditor-ready records covering agent inventory, ownership, action audit trails, and compliance evidence packaging.

— Lists record types but no access path or format. product docs

11 **Dark Mode across Agentic Fabric and ISC** NEW 45

Dark Mode support is added across SailPoint Agentic Fabric and a select number of Identity Security Cloud pages, and is user-configurable per account. Custom branding colors are not applied in Dark Mode and continue to appear only in Light Mode.

— Describes scope and branding caveat, no toggle location named.

Enhancement: Dark Mode Support Across S…

12 **Continuous shadow AI monitoring and alerts** NEW 42

Continuous post-activation monitoring fires shadow AI alerts for unsanctioned tools and updates identity correlations in real time.

— Describes alerting behaviour, no configuration detail. product docs

13 **Embedded Identity Graph per agent** NEW 40

An embedded Identity Graph is now available per agent and per application for visualizing relationships between non-human identities.

— Describes purpose only; no interaction detail. product docs

BREAKING ON UPGRADE

! Custom branding colors will not be applied in Dark Mode — they continue to appear in Light Mode only.

WAS THIS USEFUL?

◆ Data Security & DLP

Immuta

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06 NOTES ↗

Immuta is a data governance platform that controls access to sensitive data and ensures compliance through automated policies and masking.

> Immuta shipped a native Trino/Starburst integration with direct policy enforcement on Trino objects, alongside a Connections UI overhaul adding health-state indicators, per-object policy sync controls, tag inheritance, and Redshift automated data discovery — plus several breaking changes affecting Databricks query audit and Snowflake policy roles.

└──▷ WHAT SHIPPED · 12 FEATURES most completely described first ? what's the number?

01 **Native Trino/Starburst integration** NEW 70

A new Trino integration lets teams register a Trino connection so Immuta can administer access controls and enforce Immuta policies directly on objects and end-user access within Trino. It includes customizable Trino plugin properties to tailor the plugin configuration, write access mapping to customize how read/write access policies are enforced for Starburst (Trino), rotation of the system API key for the integration, and supported authentication methods and audit features for compliance.

– Describes mechanism and config surfaces but no exact commands product docs

02 **Automated data discovery for Redshift viewless integration** NEW 65

Automated sensitive data discovery now supports Immuta's Amazon Redshift viewless integration, covering regex, dictionary, and column-name-regex identifier types.

– Names integration and identifier types supported snapshot-20260806

03 **Per-object policy sync toggle** NEW 60

A new per-object policy sync setting in Connections controls whether Immuta pushes policy down to the data platform for a given data object, showing a Custom label when a child object's setting differs from its parent's.

– Names the setting and UI label, no exact steps snapshot-20260806

04 **Configurable New tag behavior** NEW 60

App Settings now includes a configurable New tag behavior, enabled by default for columns and disabled by default for data sources.

– Names exact setting location and defaults snapshot-20260806

05 **New Immuta-managed built-in groups** NEW 60

Three new Immuta-managed built-in groups — `All Agents`, `All Users`, and `All Identities` — are now available for use in subscription policies and data policies.

— Names the exact groups and where they apply `snapshot-20260806`

THINNER COVERAGE BELOW

06 **Connection health state indicator** NEW 55

Connections in the Connections UI now display an `ACTIVE`, `INACTIVE`, or `UNHEALTHY` health state for Databricks Unity Catalog and Snowflake integrations.

— Names states and UI location but no workflow detail `snapshot-20260806`

07 **Tag inheritance at database/schema level** NEW 50

Tags can now be inherited at the database or schema level, with bulk assignment or removal propagating to all attached data sources.

— Explains behavior but no config path given `snapshot-20260806`

08 **Snowflake resync covers subscription policies** IMPROVED 50

Snowflake data policy resync now resyncs both data policies and subscription policies when triggered from a data source health check.

— Names trigger and scope, no manual steps given `snapshot-20260806`

09 **Force-delete unhealthy connections** NEW 45

Connections in an `UNHEALTHY` state can now be force-deleted.

— Clear capability, minimal mechanism explained `snapshot-20260806`

10 **Integration settings moved into Connections UI** IMPROVED 45

Integration settings for Databricks Unity Catalog and Snowflake connections now live directly within the Connections UI rather than in a separate settings area.

— Names new UI location, little further detail `snapshot-20260806`

11 **Detail pages for leaf data objects** NEW 40

New detail pages exist for leaf data objects in Connections, each including a direct link to the corresponding data source.

— Names UI surface, little mechanism detail `snapshot-20260806`

12 **Snowflake integration error banners removed** BREAKING 40

Integration error banners for Snowflake integration validation failures will no longer be displayed; errors will instead appear in the new Connections experience.

— States change but not exact new error location `snapshot-20260806`

└──▷ BREAKING ON UPGRADE

! Starting August 26, 2026, Databricks will mask `query_text` in the `query_history` system table by default; Immuta query audit will no longer contain query text and all query audit events will show as `Indeterminant` unless the Immuta service principal is added to the `databricks_pii_access` group in Databricks.

! Immuta policy role names used to grant users access to Snowflake data sources will be renamed and recreated as part of the Snowflake integration architecture update rolling out from July 23, 2026 (user role names remain unchanged).

! Integration error banners for Snowflake integration validation failures will no longer be displayed; errors will instead appear in the new Connections experience.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

OpenVPN

Feedback ⓘ 1 RELEASE · 2026-08-06 NOTES ↗ CODE ↗

OpenVPN is an open source VPN daemon

> OpenVPN v2.7.6 simplifies default configuration by making `tun` the implicit device and enabling `TCP_NODELAY` by default, while capping keepalive timers at 24 hours.

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

01 **TCP_NODELAY enabled by default** BREAKING 78

The `TCP_NODELAY` socket flag is now always enabled on the local socket, whereas previously it required `--tcp-nodelay`. In server mode, `--tcp-nodelay` now specifically pushes `socket-flags TCP_NODELAY` to clients that may not have this default.

— Explains mechanism and server-side push behavior with flag names. v2.7.6

02 **Default device changed to tun** IMPROVED 70

Omitting `--dev` from an OpenVPN config now defaults to `--dev tun`, removing a previously required boilerplate option for the common tun case.

— Names exact flag and default change, no example. v2.7.6

03 **24-hour cap on ping and keepalive intervals** BREAKING 70

The `--ping` and `--keepalive` options now reject values exceeding 86400 seconds (24 hours); configs with larger values will need to be updated.

— Gives exact limit and required migration action. v2.7.6

└──▷ BREAKING ON UPGRADE

! The `--ping` and `--keepalive` options now reject values exceeding 86400 seconds (24 hours); configs with larger values will need to be updated.

! `TCP_NODELAY` is now enabled by default on all sockets; previously it was opt-in via `--tcp-nodelay`.

WAS THIS USEFUL? 👍 👎

Zscaler

Feedback

1 RELEASE · 2026-08-05

BLOG ↗

Zscaler is a cloud-based security platform providing zero-trust network access, threat prevention, and data protection for organizations.

> Zscaler shipped AI Guard integration with Anthropic's Claude inference hooks, extending Zero Trust DLP policy enforcement directly into the Claude model layer.

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

01 **AI Guard integration with Claude inference hooks** NEW

62

Zscaler AI Guard integrates with Anthropic Claude inference hooks to enforce Zero Trust data-loss-prevention policy natively at the Claude model layer, evaluating full conversation context in real time rather than individual messages. Existing Zscaler platform policies extend to Claude Enterprise touchpoints with no additional governance stack required, and existing Zscaler customers can activate the feature using their current policy configuration.

— Names product and integration point but no config key or endpoint. 2026-08-05-1bd71b85

WAS THIS USEFUL? 👍 👎

ScoutDNS

Feedback ⓘ 1 RELEASE · SEEN 2026-08-06 NOTES ↗

ScoutDNS is a DNS reconnaissance and security tool that maps domain infrastructure, identifies misconfigurations, and monitors DNS records for vulnerabilities.

> ScoutDNS shipped a ground-up console redesign this cycle, adding explicit match/subtree controls and comments for allow/block list rules, an auto-generated Zero Trust TLD allowlist, per-policy Secure Scores, a cross-tenant MSP dashboard, and dark mode.

└──▷ **WHAT SHIPPED · 9 FEATURES** most completely described first ? what's the number?

01 **Match/subtree option on allow/block list rules** NEW

65

Every custom allow/block list rule now carries a `match` or `subtree` option that makes explicit whether the rule covers a single hostname or a domain and all its subdomains; a single wildcard entry using subtree now covers both the domain and its subdomains, removing the need for two separate entries.

— Names the exact rule option but no config path shown

The all-new ScoutDNS console is live

02 **Auto-generated Zero Trust TLD allowlist** NEW

62

Zero Trust TLD filtering can now auto-generate a TLD allowlist from the last 30 days of real query traffic for any organization.

— Clear mechanism and lookback window, no UI path given

The all-new ScoutDNS console is live

THINNER COVERAGE BELOW

03 **Deployments subtab on Policies page** NEW

57

A new Deployments subtab on the Policies page shows every policy assignment across sites, profiles, and personas in one account-wide view, with inline editing.

— Names exact UI location and scope, no deeper mechanism

The all-new ScoutDNS console is live

04 **Per-policy Secure Score** NEW

53

Each policy now carries a 0–100 Secure Score with a per-item breakdown showing what to enable to raise it.

— Scale and breakdown named, no navigation path given

The all-new ScoutDNS console is live

05 **Stateful site health based on WANs and relays** IMPROVED

45

Sites now report stateful up/down health based on WANs and relays rather than recent query volume, changing how connectivity status is determined.

— Explains the before/after but no config or metric detail

The all-new ScoutDNS console is live

06 **MSP dashboard for cross-tenant view** NEW

45

A new MSP dashboard under the Organizations tab gives a cross-tenant view of all managed accounts.

— Names tab location, lacks detail on what's shown The all-new ScoutDNS console is live

07 **Comments on allow/block list entries** NEW

37

Comments can now be added to any allow or block list entry to preserve the reason a rule was created.

— Simple feature stated plainly, no further mechanism given

The all-new ScoutDNS console is live

08 **Dark mode in console** NEW

30

Dark mode is now available, selectable from the console header.

— Bare description with a UI location, nothing more The all-new ScoutDNS console is live

09 **Rebuilt help site with full-text search** IMPROVED

28

The help site was rebuilt with full-text search and current screenshots for every tab and function.

— Documentation update stated with minimal further detail

The all-new ScoutDNS console is live

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.38.0 (2026-08-06) BLOG ↗

RegScale is a compliance and regulatory management platform that automates governance, risk, and compliance (GRC) workflows for organizations.

RegScale CLI 6.38.0 adds eMASS CAC authentication support for custom CA bundles (e.g. DoD PKI) via config, env var, and CLI flag.

└──▷ TRY IT

Register an eMASS API connection trusting a DoD PKI CA bundle when running in a CI/CD pipeline where env vars are preferred over config files.

```
$ EMASS_SSL_CA_CERT=/etc/ssl/dod-root-ca.pem regscale emass_api
register
```

Persist the DoD PKI CA bundle path in the RegScale config so all eMASS commands trust it without repeating the flag.

```yaml
emass:
  ssl_ca_cert: /etc/ssl/dod-root-ca.pem
```

Supply the CA bundle inline at registration time without modifying config or environment.

```
$ regscale emass_api register --ca-cert /etc/ssl/dod-root-ca.pem
```

Adds `emass.ssl_ca_cert` config key to trust a custom CA bundle (e.g. DoD PKI) when verifying the eMASS server certificate during CAC authentication.

Adds `EMASS_SSL_CA_CERT` environment variable as an alternative to `emass.ssl_ca_cert` for supplying the custom CA bundle path.

Adds `--ca-cert` option to the `emass_api register` command to specify the custom CA bundle at registration time.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Cline

Feedback ⓘ 4 RELEASES · 2026-08-05 → 2026-08-06 NOTES ↗ CODE ↗

Autonomous coding agent as an SDK, IDE extension, or CLI assistant.

> Cline this window focused on model catalog updates and a major overhaul of reasoning-effort handling, standardizing it across all AI SDK providers including Ollama, alongside a friendlier UX for retired free-model promotions.

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

01 Unified reasoning effort control across AI SDK providers

65

BREAKING

Reasoning effort settings (effort levels and enable/disable toggles) now map to a shared path across all AI SDK providers, including Ollama, instead of going through per-provider thinking overrides; disabling reasoning is respected everywhere. Per-provider thinking overrides are replaced by the AI SDK's native reasoning setting, so existing configurations relying on those overrides will need to be updated.

– Explains mechanism and migration need but no exact config key given. cli-v3.0.51 sdk/sdk/v0.0.71 v4.1.5

THINNER COVERAGE BELOW

02 meta/muse-spark-1.2-contributor added to Cline provider NEW

50

Adds `meta/muse-spark-1.2-contributor` as a selectable model on the Cline provider, alongside a refreshed model catalog.

– Names the exact model id but no further detail on catalog changes. v4.1.6 cli-v3.0.51 sdk/sdk/v0.0.71

03 Dedicated notice for retired free model promotions NEW

40

When a retired free model is requested, Cline shows a dedicated notice with a button to pick another model, instead of a generic error with only a Retry prompt.

– Describes behavior change but no UI path or config named. v4.1.5

└──▷ BREAKING ON UPGRADE

! Per-provider thinking overrides are replaced by the AI SDK's native reasoning setting; existing configurations relying on those overrides will need to be updated.

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-06 NOTES ↗ CODE ↗

Claude Code is an agentic coding tool that lives in your terminal, understands your codebase, and helps you code faster by executing routine tasks, explaining complex code, and handling git workflows - all through natural language commands.

> Claude Code v2.1.223 adds owner-level wildcard controls for marketplace allow/block lists, a `/teleport` command for moving cloud sessions to a local machine, and an upgraded `/code-review` with effort levels and PR targeting, alongside new env vars for controlling context-window behavior.

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

01 /code-review effort levels and PR targeting IMPROVED

95

Makes `/review` an alias of `/code-review`; `/code-review` now accepts an effort level and PR number via `/code-review <level> <pr#>`, with `/code-review ultra` triggering a deep cloud review (e.g. `/code-review ultra 4321`). It also remembers the last effort level used when invoked without one, so repeat reviews don't require retyping the level.

Run a deep review of a pull request straight from the terminal using the new effort-level syntax.

```
$ /code-review ultra 4321
```

— Exact command syntax and runnable example given v2.1.223

02 **Owner wildcard marketplace controls** NEW 90

Adds `"owner/*"` wildcard entries to `strictKnownMarketplaces` and `blockedMarketplaces` managed settings, allowing or blocking all marketplace repos under a GitHub org in a single entry, e.g. `{"blockedMarketplaces": ["acme-corp/*"]}`.

Block all marketplace repos from a specific GitHub org in managed settings with a single wildcard entry instead of listing each repo.

```json
{
  "blockedMarketplaces": ["acme-corp/*"]
}
```

— Exact config keys and runnable example given v2.1.223

03 **/teleport cloud-to-local session handoff** NEW 80

Adds a `claude --teleport <session id>` flag and a `/teleport` hint shown in cloud sessions, letting a user continue a cloud session locally.

Pick up a cloud session on your local machine when you need to switch environments mid-task.

```
$ claude --teleport <session id>
```

— Exact command shown, underlying mechanism unexplained v2.1.223

04 **Expanded 1M context auto-compaction cap** IMPROVED 70

Expands `CLAUDE_CODE_DISABLE_1M_CONTEXT` to cap every Claude model with a native 1M context window to 200K via auto-compaction, rather than only a fixed model list, and adds a startup warning when auto-compaction isn't holding the session to 200K.

— Named env var and behavior, no usage example v2.1.223

05 **Opt-out of unknown model context window enforcement** NEW 70

Adds `CLAUDE_CODE_DISABLE_UNKNOWN_MODEL_WINDOW_ENFORCEMENT=1` env var to opt out of the new behavior that keeps sessions on unrecognized model IDs within the assumed context window.

— Named env var given, no example command v2.1.223

THINNER COVERAGE BELOW

06 **Warning for restricted subagent model fallback** NEW

45

Adds a warning when a workflow agent, forked skill, slash command, or resumed background agent's requested subagent model is restricted and the parent model runs instead.

— Describes behavior only, no named surface or example v2.1.223

WAS THIS USEFUL? 👍 👎

OpenCode

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗ CODE ↗

The open source coding agent.

> OpenCode adds Poolside as a supported provider with self-hosted deployment configuration, and streamlines xAI (Grok) authentication with a device-code OAuth flow alongside API key access.

└──▷ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Poolside provider support with self-hosted deployment config

NEW 95

OpenCode adds Poolside as a supported provider via its OpenAI-compatible API, with API keys configured at `platform.poolside.ai` or stored locally at `~/.secrets/poolside-key` for self-hosted deployments. Self-hosted Poolside instances are set up as a custom OpenAI-compatible provider with a custom provider ID and base URL (e.g. `https://poolside.example.com/openai/v1`), discovering model IDs via the `/openai/v1/models` endpoint, and support the `reasoning_content` field plus `interleaved` reasoning mode for thinking-capable models.

Configure a self-hosted Poolside deployment as a custom provider, storing credentials outside the config file.

json

```
{
  "id": "poolside-deployment",
  "name": "Poolside deployment",
  "baseURL": "https://poolside.example.com/openai/v1",
  "apiKey": "{file:~/.secrets/poolside-key}",
  "models": [
    {
      "id": "<served-model-id>",
      "name": "Poolside deployment model",
      "reasoning": "interleaved",
      "field": "reasoning_content"
    }
  ]
}
```

— Named endpoint, file path and full JSON config example included. product docs

02 xAI (Grok) authentication via SuperGrok OAuth and API key

IMPROVED 60

OpenCode adds two xAI authentication paths: a SuperGrok subscription via device-code OAuth, where OpenCode opens xAI's verification link with the user code pre-populated and automatically polls the token endpoint, or a pay-as-you-go API key obtained from the xAI console. In v1.18.14 this login flow was simplified to a single device-code flow to better support headless and remote environments.

— Describes mechanism and simplification but gives no exact command. v1.18.14

WAS THIS USEFUL? 👍 👎

OpenHands

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

OpenHands: AI-Driven Development

> OpenHands v1.10.0 adds activity log export, a featured automations dashboard, manifest-driven automation sub-pages, faceted skills filtering, and switches Canvas's default model to GLM 5.2.

└──▷ **WHAT SHIPPED · 5 FEATURES** most completely described first ? what's the number?

01 **Faceted skills filtering** NEW — 40

Adds faceted rail filtering to the Skills page for faster skill discovery.

– Names UI location (Skills page) but no filter facets listed v1.10.0

02 **Manifest-driven automation sub-pages** NEW — 35

Adds manifest-driven sub-pages to the automation interface, enabling structured automation workflows.

– Mentions manifest-driven mechanism but no schema or file path v1.10.0

03 **GLM 5.2 default model for Canvas** IMPROVED — 35

Sets GLM 5.2 as the default model for Canvas.

– Names specific model and feature but no config or override info v1.10.0

04 **Featured automations landing dashboard** NEW — 30

Ships a featured automations landing dashboard for discovering and launching automations.

– Describes purpose but no navigation path or details v1.10.0

05 **Activity log export** NEW — 25

Adds activity log export capability to the interface.

– Names the capability but no format or location given v1.10.0

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

Framework for orchestrating role-playing, autonomous AI agents. By fostering collaborative intelligence, CrewAI empowers agents to work together seamlessly, tackling complex tasks.

> CrewAI 1.15.12 consolidates project scaffolding into a single command, adds a URL-reading tool for agent workflows, and attaches app metadata to platform action tools.

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

01 Unified `crewai create <resource>` scaffolding IMPROVED

70

Replaces separate scaffold commands with a single unified `crewai create <resource>` command for generating new CrewAI resources such as crews or flows.

Scaffold a new CrewAI resource (crew, flow, etc.) using the unified create command.

```
$ crewai create <resource>
```

— Names exact command, runnable as given 1.15.12

THINNER COVERAGE BELOW

02 URLReadTool for reading arbitrary URLs NEW

40

Adds `URLReadTool`, letting agent workflows read arbitrary URLs.

— Names the tool but no usage detail 1.15.12

03 App metadata on platform action tools IMPROVED

20

Adds app metadata to platform action tools, per the release notes.

— Bare mention, no mechanism or fields named 1.15.12

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama

Feedback ⓘ 1 RELEASE · 2026-08-04 NOTES ↗ CODE ↗

Get up and running with Kimi-K2.6, GLM-5.2, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

> Ollama v0.32.6 brings OpenAI-compliant streaming usage reporting and automatic speculative decoding for Qwen3.5 on Apple GPUs, while removing the experimental image generation feature.

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

01 stream_options.include_usage in chat completions API

IMPROVED

90

Adds `stream_options.include_usage` support in `/v1/chat/completions` so usage is returned in a separate trailing chunk, matching OpenAI's streaming wire format — `role` now appears only on the first chunk and `finish_reason` on its own chunk.

Stream a chat completion and receive token usage in a separate trailing chunk, matching OpenAI's wire format.

```
$ curl http://localhost:11434/v1/chat/completions -d '{"model": "gemma4", "stream": true, "stream_options": {"include_usage": true}, "messages": [{"role": "user", "content": "Why is the sky blue?"}]}'
```

— Named endpoint, field, and runnable curl example provided v0.32.6

THINNER COVERAGE BELOW

02 Automatic MTP speculative decoding for Qwen3.5 on Apple GPUs

IMPROVED

50

Qwen3.5 inference on Apple GPUs is faster because the MLX engine now automatically uses the model's MTP head for speculative decoding.

— Explains mechanism but no config or flag to invoke it v0.32.6

03 Experimental image generation removed

BREAKING

45

Experimental image generation has been removed in v0.32.6; users must continue using v0.32.5 for image generation support.

— Clear removal noted with version to revert to v0.32.6

└──▷ **BREAKING ON UPGRADE**

! Experimental image generation is no longer available in v0.32.6; the feature has been removed and requires staying on v0.32.5.

WAS THIS USEFUL?  

GOVERN

◆ AI/LLM Security

ToolHive

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive's v0.42.0 release introduces a full CLI and REST API for AI-tool plugin management via `thv ai-plugin`, complete with a registry catalog and Sigstore signature verification, alongside concurrent stdio-server sharing across MCP clients — and ships several breaking changes touching CRD status fields, Cedar authorization ordering, panic logging, and internal telemetry APIs.

└──▷ WHAT SHIPPED · 8 FEATURES most completely described first ? what's the number?

01 **Removed status.referencingWorkloads fields from config CRDs**

90

BREAKING

`status.referencingWorkloads` and `status.referenceCount` fields (and the References printer column) are removed from all six config CRDs (`MCPOIDCConfig`, `MCPAuthzConfig`, `MCPExternalAuthConfig`, `MCPToolConfig`, `MCPWebhookConfig`, `MCPTelemetryConfig`); replace any automation that reads them with workload field queries.

— Names exact CRDs, fields, and required migration path v0.42.0

02 **Cedar authorization evaluated post-mutation** BREAKING

90

Cedar authorization policy is now evaluated against the post-mutation MCP request — if you run a mutating webhook (`--webhook-config` with a `mutating:` entry, or `MCPWebhookConfig.spec.mutating`) alongside Cedar authz, policy decisions and audit records (`type`, `target.name`) can change on upgrade; re-audit policies against the post-mutation shape before upgrading.

— Details exact flags, audit fields, and required re-audit action v0.42.0

THINNER COVERAGE BELOW

03 **thv ai-plugin CLI and REST API** NEW

55

Adds `thv ai-plugin` subcommand with a full CLI and REST API for end-to-end AI-tool plugin management.

— Names command and API but no endpoints or flags listed v0.42.0

04 **Removed pkg/telemetry/providers and optimizerdec constants** 55

BREAKING

`pkg/telemetry/providers` was deleted and two long-published `optimizerdec` constants were removed from the Go API, breaking any out-of-tree importers that reference them.

— Names removed package and constants but no migration guidance v0.42.0

05 **Sigstore signature verification for plugins** NEW 50

Adds Sigstore signature verification for skills/plugins at install, sync, and upgrade time.

— Names trigger points but no config flag or command v0.42.0

06 **Recovered HTTP panics no longer logged** BREAKING 50

Recovered HTTP panics no longer emit a `slog.Error` log line or stack trace — setups relying on log-based alerts or support bundles to diagnose recovered panics will silently lose that diagnostic output (a patch fix is expected).

— Explains impact and notes upcoming fix, no interim action given v0.42.0

07 **ai-plugin registry catalog** NEW 40

Adds a registry catalog to `thv ai-plugin` for discovering and deploying plugins from a curated index.

— Describes purpose but no catalog commands or fields given v0.42.0

08 **Concurrent stdio server sharing across MCP clients** IMPROVED 30

Enables multiple MCP clients to share a single stdio server concurrently.

— Brief statement with no mechanism or configuration detail v0.42.0

└──▷ BREAKING ON UPGRADE

! `status.referencingWorkloads` and `status.referenceCount` fields (and the References printer column) are removed from all six config CRDs (`MCPOIDCConfig` , `MCPAuthzConfig` , `MCPExternalAuthConfig` , `MCPToolConfig` , `MCPWebhookConfig` , `MCPTelemetryConfig`); replace any automation that reads them with workload field queries.

! Cedar authorization policy is now evaluated against the post-mutation MCP request — if you run a mutating webhook (`--webhook-config` with a `mutating:` entry, or `MCPWebhookConfig.spec.mutating`) alongside Cedar authz, policy decisions and audit records (`type` , `target.name`) can change on upgrade; re-audit policies against the post-mutation shape before upgrading.

! Recovered HTTP panics no longer emit a `slog.Error` log line or stack trace — setups relying on log-based alerts or support bundles to diagnose recovered panics will silently lose that

diagnostic output (a patch fix is expected).

! `pkg/telemetry/providers` was deleted and two long-published `optimizerdec` constants were removed from the Go API, breaking any out-of-tree importers that reference them.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 6, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue