THE CYBER TOOLCHAIN

NO. 152

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 28, 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

INFORMATION DENSITY

everything



BRIEFED | **EVERYTHING**

Every tool in the issue, in full. all 172 changes

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 47 of 51 tools

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: **GPT-5.5** Opus 5

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out today: StepSecurity can now stop fresh GitHub Action versions before they enter CI, Hayabusa can read

vim ? help

Windows event logs straight from forensic images, and Veracode lowers the friction of recording authenticated DAST crawls.

PROTECT

Hold back newly published GitHub Actions before CI trusts them

The new GitHub Action Cooldown check applies minimum-age policy to Actions and reusable workflows, not just packages. That gives teams a buffer against compromised or maliciously published Action versions without manually policing every workflow update.

StepSecurity

RESPOND

Analyze Windows event logs from disk images without mounting them first

Hayabusa can now scan event logs directly from VHDX and EWF forensic images across its event-log commands. Investigators can run detection over acquired media without needing admin rights, extraction, or a mount workflow that can slow triage and complicate evidence handling.

Hayabusa

PROTECT

Record authenticated web scans without a separate Selenium workflow

The experimental `veracode dynamic record` command captures DAST login and crawl scripts from the CLI. That removes a common handoff step for authenticated scanning, where teams otherwise have to maintain browser-recorder artifacts outside the scanner workflow.

Veracode

IDENTIFY

Find overly broad cloud and SaaS exposure in more places

Prowler added sovereign/custom Azure and M365 regions plus a check for permissive GitHub Actions OIDC trust in AWS IAM roles, while Zanshin expanded leaked-credential and AI-posture coverage and Plerion added Azure network checks. The practical gain is broader posture inventory across environments that are often missed or reviewed manually.

Prowler · Zanshin · Plerion

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Password & Credential Cracking

brutespray

Feedback ⓘ v2.7.9 NOTES ↗ RANK

BruteSpray is a Python tool that takes Nmap GNMAP output and automatically performs credential brute-forcing against discovered services using Medusa.

> BruteSpray v2.7.9 adds a round-robin credential scheduling mode for lockout-aware brute-forcing and reworks host pooling so services sharing a host:port are tracked independently.

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

01 **Round-robin credential scheduling** Not checked yet NEW

87

A new `-schedule round-robin` option interleaves credentials across all active targets, sending one credential per host per rotation to help avoid account lockouts. It requires `-T >= 2`, and adding `-t 1` gives a deterministic per-host order.

Spray a list of targets one credential at a time per host per round to avoid account lockouts, with one thread per host for a predictable rotation order.

```
$ brutespray -f nmap.gnmap -u users.txt -p passwords.txt -
schedule round-robin -T 4 -t 1
```

▶ View command output v2.7.9

— Names exact flags, requirement, and a runnable example command. v2.7.9

THINNER COVERAGE BELOW

02 **Per-service host pool keying** Not checked yet IMPROVED 52

Host pools are now keyed by `host:port/service` so that two different services running on the same `host:port` each get an independent job queue instead of sharing one.

— Explains mechanism but gives no direct command or flag to invoke it. v2.7.9

WAS THIS USEFUL? 👍 👎

◆ Autonomous Offensive Frameworks (AOF)

xalgorix

Feedback ⓘ v4.6.212 NOTES ↗ RANK

> xalgorix added an environment variable to reduce per-request token usage on long scans by sending a compact system prompt instead of the full methodology on every call.

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

01 Compact core prompt via environment variable

Not checked yet NEW

85

Setting `XALGORIX_SYSTEM_PROMPT_CORE_ONLY` sends a compact core prompt on subsequent requests instead of re-sending the full 22-phase methodology, payload libraries, and playbooks on every call, reducing per-request token usage on long scans. The setting defaults to off and can be persisted in `~/.xalgorix.env`.

Persist the setting so every xalgorix run on this host uses the compact core prompt without re-exporting the variable each time.

ini

```
XALGORIX_SYSTEM_PROMPT_CORE_ONLY=true
```

— Names exact env var, default, and config file location. v4.6.212

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Self-Hosted

Feedback ⓘ

changelog · Nov 28, 2026 | CHANGELOG ↗ | RANK

GitGuardian Self-Hosted is a secrets detection platform deployed on-premises that scans source code and git repositories to prevent sensitive data leaks.

> GitGuardian Self-Hosted hardened its Redis-backed credential cache with encryption at rest and updated its bundled ggscout dependency.

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

01 **ggscout 0.36.1 bundled in Helm chart** Not checked yet IMPROVED

55

GitGuardian Self-Hosted now bundles `ggscout 0.36.1` in the Helm chart; validity and permissions checks require `ggscout 0.34.1` or later on any out-of-cluster instances.

— Names exact version requirement, a clear upgrade action changelog-20261128-4370ae17

02 **Encrypted Redis cache for credentials and tokens** Not checked yet

IMPROVED

30

Integration credentials and access tokens stored in GitGuardian's Redis cache are now encrypted at rest.

.
.
.

— Describes what changed but no mechanism or scope detail changelog-20261128-4370ae17

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.167.17 NOTES ↗ RANK

Infisical is an open-source secret management platform that helps teams securely store, sync, and manage API keys, environment variables, and other sensitive credentials.

> Infisical v0.167.17 introduces a remote MCP server for machine-to-machine protocol access, Azure managed identity authentication across all Azure connection types, and a DNS Server (RFC 2136) ACME DNS-01 challenge provider in PKI, alongside PAM session limits, sync fixes, and several smaller API/UI additions.

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

01 **Helm chart instance setup token** Not checked yet NEW 65

Helm chart bootstrap now accepts `instance_setup_token` as a configuration value.

— Exact config key named, runnable for Helm deployments v0.167.17

02 **DNS Server ACME challenge provider in PKI** Not checked yet NEW 60

Infisical's PKI now supports DNS Server (RFC 2136) as an ACME DNS-01 challenge provider for certificate issuance.

— Names protocol/RFC but no setup steps v0.167.17

03 **Import endpoint returns project ID** Not checked yet IMPROVED 60

The `POST` import endpoint now returns `importEnv.projectId` in its response body.

— Exact endpoint and response field named v0.167.17

THINNER COVERAGE BELOW

04 **Folder structure preserved in AWS secret syncs** Not checked yet IMPROVED 55

Secret syncs to AWS Parameter Store and Secrets Manager now preserve folder structure during recursive syncs.

— Names two AWS targets and the sync behavior fixed v0.167.17

05 **Azure managed identity for app connections** Not checked yet NEW 50

Azure app connections now support managed identity authentication across all Azure connection types.

– Names scope (all connection types) but no config steps v0.167.17

06 **Maximum access duration in Privileged Access Management**

Not checked yet NEW 40

Privileged Access Management now enforces a maximum access duration on sessions.

– States behavior, no limit value or config key given v0.167.17

07 **Custom role creation from add-members dialog** Not checked yet

IMPROVED 40

Create a custom role directly from the add-members dialog without leaving the workflow.

– UI path given but no field or permission detail v0.167.17

08 **Custom SAML MFA attribute values** Not checked yet IMPROVED 35

Infisical now accepts custom MFA attribute values for SAML authentication.

– No attribute names or config path specified v0.167.17

09 **Delete revoked honey tokens from UI** Not checked yet IMPROVED 35

Revoked honey tokens can now be deleted from the UI.

– Clear UI action, minimal further detail v0.167.17

10 **Remote MCP server** Not checked yet NEW 25

Infisical adds a remote MCP server for machine-to-machine protocol access.

– No mechanism or setup detail given v0.167.17

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS Feedback ⓘ

Add collaborators to secret incidents, alongside the assignee (2026-11-27)

BLOG ↗ RANK

GitGuardian SaaS is a secrets detection platform that scans code repositories and developer workflows to identify and remediate exposed API keys, credentials, and sensitive data.

GitGuardian adds multi-person collaborators to secret incidents, with bulk editing, incident-list filters, CSV export column, API access, and a new change event for existing integrations.

Add collaborators to secret incidents in Internal Monitoring and Public Monitoring from the new Collaborators field in the incident side panel, beneath Assignee. Not checked yet

Add or remove collaborators in bulk from the incident list. Not checked yet

Filter incident lists by Collaborators or Has collaborators, including in saved views and AI filters, to find incidents any teammate is working on. Not checked yet

The incident CSV export now gains a `collaborators` column. Not checked yet

A new Collaborators change event is now available for Slack, Microsoft Teams, Jira, ServiceNow, and webhook integrations. Not checked yet

+1 more

UPDATED INTEGRATION Confluence Cloud Integration — scan previous page and blog post versions

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

caido

Feedback ⓘ v0.60.0 NOTES ↗ RANK

Caido is a lightweight web security auditing tool that acts as an HTTP proxy, allowing users to intercept, inspect, and manipulate web traffic.

> Caido v0.60.0 adds native out-of-band interaction detection, environment variable support in match/replace rules, and workspace settings import/export, alongside project compression and improved theming.

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

01 **Environment variables in match/replace rules** Not checked yet NEW

50

Use environment variables inside match and replace rules to inject dynamic values without hardcoding secrets.

– Explains mechanism but no concrete syntax example v0.60.0

02 **Native out-of-band (OOB) integration** Not checked yet NEW

40

Caido adds a native out-of-band (OOB) integration for detecting server-side callbacks during testing.

– Describes purpose but no setup or config details v0.60.0

03 **Import/export shortcut and sidebar settings** Not checked yet NEW

40

Import and export shortcut and sidebar settings to share or back up your workspace configuration.

– Clear feature but no file format or UI path named v0.60.0

04 **Project storage compression** Not checked yet IMPROVED

20

New projects now use compression to reduce storage footprint.

– Bare statement without mechanism or figures v0.60.0

05 **Light theme and custom theme support** Not checked yet IMPROVED

20

Caido adds a light theme and better custom theme support for the UI.

– Cosmetic change with no further detail v0.60.0

WAS THIS USEFUL? 👍 👎

Acunetix 360 Feedback ⓘ v26.11.1 – 26 November 2026 (2026-11-26)

BLOG ↗

Acunetix 360 scans web applications and APIs to identify security vulnerabilities through automated testing.

NEW VULNERABILITY CHECK New security check

WAS THIS USEFUL? 👍 👎

◆ Code Security

syzkaller

Feedback ⓘ changes since 2026-11-25 CODE ↗

Syzkaller is an unsupervised coverage-guided kernel fuzzer developed by Google that automatically finds bugs in operating system kernels by generating and executing random syscalls.

syzkaller adds console-only crash scanning, shadow VMCS and Hyper-V eVMCS support, and path-based subsystem extraction

Enable console-only crash scanning for raw C runs with WithConsoleOnlyCrash(), so only crashes visible on the real console (not stdout/stderr) are reported. Not checked yet

syzkaller now supports shadow VMCS in x86 SYZOS via new `SYZOS_API_NESTED_INTEL_VMREAD` (346), `SYZOS_API_NESTED_INTEL_VMWRITE` (347), `NESTED_PAGE_VMREAD_BITMAP` (70), `NESTED_PAGE_VMWRITE_BITMAP` (80), and `NESTED_PAGE_SHADOW_VMCS` (90) API surfaces. Not checked yet

syzkaller now supports Hyper-V eVMCS in x86 SYZOS via `SYZOS_API_NESTED_INTEL_EVMCS_ENABLE` (348) and `SYZOS_API_NESTED_INTEL_EVMCS_WRITE_MASK` (349), with `nested_evmcs_flags` and `NESTED_EVMCS_FIELD_*` controls for enlightened VMCS fields. Not checked yet

Adds path-based subsystem extraction in `pkg/subsystem` to identify kernel subsystems from file paths. Not checked yet

Dashboard now uses AI root cause files for subsystem detection in addition to path-based inference. Not checked yet

WAS THIS USEFUL? 👍 👎

cve-lite-cli

Feedback ⓘ v1.44.0 NOTES ↗ RANK

> cve-lite-cli v1.44.0 adds a `packages` field to JSON diagnostics and hardens lockfile and vulnerability-lookup handling across npm, pnpm, Yarn Berry and Bun, with several changes that alter scan results or exit codes.

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

01 Per-endpoint OSV result caching

Not checked yet BREAKING

75

OSV cache entries are now kept per `--osv-url` endpoint, so scans against a custom endpoint no longer reuse results cached from the default one. The cache format changed with this release, so the first scan after upgrading runs with a cold OSV cache.

Scan a project against a private OSV mirror; each endpoint's results are cached independently so switching back to the default never pulls stale private data.

```
$ cve-lite /path/to/project --osv-url
https://osv.internal.example.com --fail-on high
```

— Names flag, cache behaviour, and includes runnable example v1.44.0

02 `--prod-only` fix for pnpm v9 lockfiles

Not checked yet BREAKING

75

`--prod-only` on a pnpm v9 lockfile now correctly excludes packages reached only through `devDependencies`, reporting only production findings; as a result these scans report fewer findings than before the fix.

Scan only production dependencies in a pnpm v9 workspace to keep CI noise focused on what ships to users.

```
$ cve-lite /path/to/project --prod-only --fail-on medium
```

— Exact flag, lockfile version, before/after, and example command v1.44.0

03 Partial-scan marking on npm advisory lookup failure

Not checked yet BREAKING

70

When the npm bulk advisory lookup fails or returns unreadable data, the scan is now marked partial instead of clean, and exits 3 under `--incomplete-policy error`.

— Exact exit code and policy flag named `v1.44.0`

04 **Stricter Yarn Berry and Bun lockfile parsing** Not checked yet

65

BREAKING

A Yarn Berry `yarn.lock` that is not valid YAML or repeats an entry now stops the scan with an error instead of being partially read. A `bun.lock` that repeats a key in the same object now also stops the scan with an error instead of being partially read.

— Names both lockfile formats and exact failure conditions `v1.44.0`

THINNER COVERAGE BELOW

05 **Lockfile entries scanned by resolved package** Not checked yet

IMPROVED

45

A lockfile entry is now scanned as the package its registry tarball URL actually installs, regardless of what name or version label the entry carries.

— Mechanism described but no example or scope detail `v1.44.0`

06 **`packages` field in JSON diagnostics** Not checked yet NEW

45

JSON diagnostic output now includes a `packages` field naming which packages a given diagnostic is about.

— Names the field but no example of its use `v1.44.0`

└──▷ BREAKING ON UPGRADE

! The first scan after upgrading runs with a cold OSV cache because the cache format changed; cached results are now stored per `--osv-url` endpoint.

! `--prod-only` on pnpm v9 lockfiles now excludes `devDependencies` `-only` packages, so these scans report fewer findings than before.

! A Yarn Berry `yarn.lock` that is not valid YAML or repeats an entry now stops the scan with an error instead of being partially read.

! A `bun.lock` that repeats a key in the same object now stops the scan with an error instead of being partially read.

! Under `--incomplete-policy error`, a failed or unreadable npm bulk advisory lookup now exits 3 instead of reporting the scan as clean.

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Nov 26, 2026 BLOG ↗ RANK

The official changelog for the Snyk AI Trust Platform. Discover what's new in our security intelligence, AI-powered tools, and latest product enhancements.

> Snyk API & Web now lets you onboard API targets by pointing to a source code repository instead of supplying a schema file, with optional AI-assisted schema generation.

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

01 **API target creation from source code repository** Not checked yet

NEW 75

When adding an API target, Snyk now offers a `Repository` option in place of supplying a schema URL or file; it searches the repository for existing OpenAPI/Swagger files to generate candidate schemas. If AI usage is enabled under `Settings > Integrations > Snyk`, Snyk can additionally use AI to generate candidate schemas from source code, emailing the user when schemas are ready to review, download, and link to begin scanning.

Onboard an API target without a hand-authored spec by selecting the Repository source in the Snyk API & Web UI, so Snyk finds or generates a candidate schema for you.

📍 In Snyk API & Web, go to your API targets and choose `Add target`. When prompted for a schema source, select `Repository` instead of `URL or file`. Ensure your Snyk integration is configured at Settings › Integrations › Snyk, then pick your repository. Snyk will search for existing OpenAPI/Swagger files and, if AI usage is enabled, generate candidate schemas. You will receive an email when schemas are ready to review, download, and link to begin scanning.

— Names exact UI option and settings path but no API/endpoint detail 2026-11-26-ea96cf2f

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Veracode SCA (2026-11-26) BLOG ↗ RANK

Veracode is a cloud-based application security platform that scans code for vulnerabilities using static, dynamic, and software composition analysis testing methods.

> Veracode's biggest addition this window is an experimental `veracode dynamic record` CLI command that captures DAST login and crawl scripts without Selenium IDE, alongside a wave of SCA agent updates (malicious package detection, Poetry and Python 3.12 support, SBOM generation, Unified Policy for agent workspaces, and new `--appname` / `--show-cli` / `--json` flags), a Static Analysis Engine refresh with broader language coverage, Java/C# API wrapper enhancements, and a breaking migration of SCA agent backend URLs.

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

01 **`veracode dynamic record` CLI command for DAST authentication**

Not checked yet NEW 97

Experimental CLI command that opens Chrome, captures a login sequence or crawl script, and saves it as a SIDE/JSON file for upload to Dynamic Analysis as a login or crawl script — no Selenium IDE or browser extension required. Supports `--url`, `--output` / `-o`, `--type login` (default) or `--type crawl` (crawl scripts preserve cookies and run inside the scan's authenticated session), `--timeout` (minutes; `0` means no limit), and `--force` to overwrite existing output files. The recorder automatically inserts `deleteAllCookies`, `open`, `waitForElementPresent`, and optionally `waitForTextPresent` commands, identifying elements by `id`, `data-testid`, `aria-label`, or `placeholder` before falling back to a CSS path; TOTP-based MFA is supported by configuring a TOTP seed as a scanner variable and inserting it into the TOTP field step manually (other MFA types are not supported).

Record a login SIDE file for a web app that requires authentication before dynamic scanning, then review and replace credentials with scanner variables before upload.

```
$ ./veracode dynamic record --url https://app.example.com/login
--output login.side
```

Record a crawl script to guide the dynamic scanner through an authenticated application's surface area.

```
$ ./veracode dynamic record --url https://app.example.com/ --
output crawl.side --type crawl
```

Record a login script for a Dynamic Analysis scan using a test account, then save it as a SIDE file ready for upload.

```
$ veracode dynamic record --type login
```

Record a crawl script that runs in the scan's authenticated session (cookies are preserved, not cleared).

```
$ veracode dynamic record --type crawl
```

Capture a login sequence for a web app so dynamic scans can authenticate — no Selenium IDE or browser extension needed.

```
$ veracode dynamic record
```

Record a login sequence for a web app and save it as the default `login.side` file, ready to upload to DAST as a login script.

```
$ ./veracode dynamic record --url https://app.example.com/login
```

Re-record a login sequence with a 30-minute time limit, overwriting an existing output file so CI pipelines can refresh scripts without manual cleanup.

```
$ ./veracode dynamic record --url https://app.example.com/login
-o login.side --force --timeout 30
```

Capture a login sequence or crawl script as a SIDE file to feed into a Veracode dynamic analysis scan.

```
$ veracode dynamic record
```

Record a DAST login sequence in Chrome and save it as a SIDE file for upload as a login script — without installing Selenium IDE.

```
$ veracode dynamic record
```

— Full flag set and underlying mechanism documented across many releases. product docs

02 **Java and C# API wrapper enhancements** Not checked yet NEW

95

The Java API wrapper (v26.10.20.0) now automatically includes Git repository and commit metadata (branch, commit, author, repository, and Git provider) in Static Analysis scans submitted from a Git repo for JAR, WAR, EAR, and ZIP archives, and the `-repourl` parameter (v26.10.20.1) controls whether it updates the repository URL in the application profile. It adds page size support when calling `/appsec/v1/applications` (v26.3.19.0) for efficient retrieval of large application lists, support for the paginated Sandbox REST API, the Findings REST API to access open and mitigated findings, and a retry mechanism for the Applications API (`appsec/v1/applications`) to handle transient failures, timeouts, and intermittent network issues (v26.3.19.0, v26.2.18.0). It also adds `repourl` to update the Git repo URL, `failbuild` to fail the build if any packaged artifact fails to upload, and `publishedscansonly` to filter scans (v24.10.15.0, v24.7.14.0, v24.4.13.0). The C# API wrapper adds the same Git repo URL update support, a `failbuild` parameter, and an `-includenewmodules` parameter (v25.2.10.2, v25.2.10.2, v24.1.10.1).

Override the Git repository URL recorded in the Veracode application profile during a scan submission, useful when your CI runner checks out via a different remote URL than the canonical one.

```
$ java -jar VeracodeJavaAPI.jar -vid <api_id> -vkey <api_key> -action uploadfile -filepath target/myapp.jar -repourl https://github.com/org/repo
```

Fail a CI build immediately if any packaged artifact fails to upload, preventing a silent partial scan from proceeding.

```
$ java -jar VeracodeJavaAPI.jar -vid <api_id> -vkey <api_key> -action uploadfile -filepath target/myapp.jar -failbuild
```

— Every wrapper parameter and API endpoint named with exact versions. product docs

03 **SCA agent backend and download URL migration** Not checked yet

85

BREAKING

The SCA agent download URL changed from `https://download.sourceclear.com` to `https://sca-downloads.veracode.com` ; the Commercial region backend API URL changed from `https://api.sourceclear.io` to `https://sca-api.veracode.com` ; the European region backend API URL changed from `https://platform-backend.analysiscenter.veracode.eu` to `https://sca-api.veracode.eu` ; and the

US Federal region backend API URL changed from `https://platform-backend.analysiscenter.veracode.us` to `https://sca-api.veracode.us`. Legacy URLs remain available for 12 months before retirement.

— All four URLs and the 12-month migration window named verbatim.

Veracode SCA (2026-11-26)

---

04 **SCA Agent-Based Scan Enhancements** Not checked yet NEW

75

The SCA agent scan command adds `--appname` to link agent scan results to an application profile (Commercial region; requires OAuth or HMAC authentication) and `--show-cli` combined with `--json` to save structured JSON output to a file while still displaying a human-readable summary in the terminal; SCA upload and agent-based scans now automatically detect and report malicious packages with no extra configuration or agent download required; Unified Policy can now be assigned to SCA agent-based scan workspaces in the US Federal and European regions to create issues and break builds based on defined criteria; and SCA agent-based and SCA Upload scans add support for Python projects using the Poetry package manager, SCA agent-based scans support Python 3.12 environments, and SBOM files can be generated during SCA agent-based scans.

— Flags named but no example command showing invocation. Describes behavior but no config surface or output detail. Names regions and policy feature but no config steps given. Three thin additions grouped; no commands or config given.

Veracode SCA (2026-11-26)

---

05 **Static Analysis Engine v2026.11 language and detection updates**

Not checked yet IMPROVED

65

Adds support for Angular 21, NestJS 11.1.x, Node.js 24.x, and Scala 3.10. Improves flaw detection, including false-positive reduction, for `CWE-259`, `CWE-798`, `CWE-915`, `CWE-252`, `CWE-326`, `CWE-757`, and `CWE-311`, improves detection in code using `sqlite3` and third-party code associated with `react-native-keychain`, and improves JSP file processing and Java source code scanning.

— Concrete CWE and framework list but no user action required. product docs

THINNER COVERAGE BELOW

06 **Crawl scope and allowed-host controls for web app scans**

Not checked yet NEW 55

Web application scan configuration now lets you set crawl depth via directory-only, directory-and-subdirectories, or full parent-and-subdirectory scope, and allow crawling of both HTTP and HTTPS versions of a URL via "Add http:// and https:// versions of this URL to allowed hosts". Directory restrictions can also be set when adding an allowed URL to a DAST web application target.

— UI options named but no exact navigation path given. product docs

07 **Deep-link redirect to SCA results after login** Not checked yet

IMPROVED 20

Deep-link support now redirects unauthenticated users to login and back to the original SCA results page.

— Single-sentence behavior change with no mechanism detail. Veracode SCA (2026-11-26)

▷ BREAKING ON UPGRADE

! The SCA agent backend API URL for European region customers changed from `https://platform-backend.analysiscenter.veracode.eu` to `https://sca-api.veracode.eu`; legacy URL remains available for 12 months.

! The SCA agent backend API URL for US Federal region customers changed from `https://platform-backend.analysiscenter.veracode.us` to `https://sca-api.veracode.us`; legacy URL remains available for 12 months.

! The Commercial region SCA agent backend API URL changed from `https://api.sourceclear.io` to `https://sca-api.veracode.com`; the legacy SourceClear URL will be retired after 12 months.

! The SCA agent download URL changed from `https://download.sourceclear.com` to `https://sca-downloads.veracode.com`; the old URL will stop working after 12 months.

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

fcaptcha

Feedback ⓘ v1.53.0 NOTES ↗

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/webdecoy/fcaptcha:1.53.0
```

git

```
$ git clone --branch v1.53.0 https://github.com/WebDecoy/FCaptcha.git
# already have the repo? check out this version:
$ git checkout v1.53.0
```

published under the project's GitHub account · checked 28 Nov Wrong command?

NEW INTEGRATION Nuxt 4 module (@webdecoy/fcaptcha-nuxt)

NEW INTEGRATION Django 5.2 LTS and Laravel 13 form integrations

NEW INTEGRATION Authentik 2026.8.7 Captcha Stage recipe

NEW SUPPORTED PLATFORM Web Bot Auth verification for Python server

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Cloudsmith

Feedback ⓘ

CLI v1.32.0: Usage telemetry, with separate controls for usage and errors (2026-11-26)

BLOG ↗ RANK

Cloudsmith is a fully-managed artifact management, software supply chain security, and software distribution platform.

> Cloudsmith's CLI v1.32.0 introduces pseudonymous usage telemetry enabled by default, tracking command usage and error rates across environments including CI.

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

01 **Pseudonymous CLI usage telemetry** Not checked yet NEW 37

The Cloudsmith CLI now collects pseudonymous usage telemetry covering which commands, options, and setups are used and how often commands fail. This is enabled by default, including in CI environments.

— Describes what's tracked but no opt-out flag or config key named

CLI v1.32.0: Usage telemetry, with sepa…

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ

GitHub Checks: GitHub Action Cooldown (2026-11-26) BLOG ↗ RANK

StepSecurity secures GitHub Actions CI/CD pipelines by hardening workflows, detecting anomalies, and preventing supply chain attacks through automated security controls.

> StepSecurity introduced a GitHub Action Cooldown check that extends its existing Package Cooldown protection to GitHub Actions and reusable workflows, blocking pull requests from adopting newly published Action versions until a minimum age is met.

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

01 **GitHub Action Cooldown check** Not checked yet NEW 87

A new GitHub Action Cooldown control under GitHub Checks > Configuration can be enabled as a Required or Optional check, applying the same waiting-period protection that Package Cooldown already applies to npm, PyPI, Maven, and NuGet dependencies, but for GitHub Actions and reusable workflows. Admins set a minimum number of days an Action version must be published before a pull request can reference it, and the check evaluates only new or changed Action references in that pull request. Specific owners (`example-org/*`), any version of a specific Action (`example-org/example-action`), or a single pinned version (`example-org/example-action@v1`) can be exempted from the cooldown requirement.

Turn on the GitHub Action Cooldown gate in your GitHub Checks configuration so newly published Actions must age at least 7 days before any pull request in your org can adopt them.

> 📍 In the StepSecurity console, go to GitHub Checks › Configuration, enable GitHub Action Cooldown, set the Cooldown Period to your desired number of days (e.g. 7), and save as a Required check.

— Names config path, exemption syntax, and setup steps precisely.

GitHub Checks: GitHub Action Cooldown (…

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ

iOS Forensic Toolkit 10.14 adds agent-based file system extraction for current-generat

BLOG ↗ RANK

ElcomSoft is a try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

> ElcomSoft's iOS Forensic Toolkit gained over-the-air sideloading for its low-level extraction agent, letting examiners deploy it even on devices with damaged ports.

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

01 **OTA sideloading of low-level extraction agent** Not checked yet

NEW SUPPORTED PLATFORM NEW 45

iOS Forensic Toolkit now supports over-the-air sideloading of its low-level extraction agent, even when ports are damaged, removing reliance on a working physical port for deployment.

— Describes capability but no exact command or workflow steps given. product docs

> iOS Forensic Toolkit 10.14 adds agent-based file system extraction for current-gener
>
> BLOG ↗ LATEST STABLE

WAS THIS USEFUL? 👍 👎

Yamato Security Hayabusa

Feedback ⓘ v4.5.0 NOTES ↗

CODE ↗ RANK

Hayabusa is a fast Windows event log threat hunting and forensic timeline generator written in Rust, using Sigma-based detection rules.

> Hayabusa now scans Windows event logs directly from VHDX and EWF (E01/Ex01) forensic disk images—without mounting, admin rights, or extraction—across all its event-log commands, and v4.5.0 shrinks the initial `update-rules` download from ~37 MB to ~8 MB via a shallow clone.

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

01 **Direct VHDX and E01/Ex01 forensic image scanning** Not checked yet

NEW 98

All event-log analysis commands (`dfir-timeline` , `log-metrics` , `csv-timeline` , `scan` , `config`) now accept the `--vhdx` and `--e01` flags combined with `-d, --directory <DIR>` or `-f, --file <FILE>` . `--vhdx` reads NTFS volumes inside VHD/VHDX images on Linux, macOS and Windows with no FUSE mount, mount driver, administrator rights or temporary extraction; directory scans recurse three subdirectory levels deep and evidence paths take the form `host.vhdx!/path/log.evtx` ; BitLocker-encrypted images and images requiring an external parent are not supported. `--e01` reads physical EWF1 (`.E01`) and EWF2 (`.Ex01`) images by pointing at the first segment while keeping remaining segments in the same directory, preserving evidence paths like `disk.E01!/Windows/System32/winevt/Logs/Security.evtx` ; logical `.L01` / `.Lx01` , password-protected EWF and BitLocker images are not supported, and the feature requires Rust 1.99 or newer. `--vhdx` and `--e01` can be combined to scan a mixed directory in one pass, `-J` can be combined with either to pull JSON-format logs from inside an image, and partial/incomplete scans still produce readable HTML reports and pivot files, are marked incomplete, and exit nonzero.

Build a DFIR timeline from a forensic E01 disk image without first extracting its event logs.

```
$ hayabusa dfir-timeline -f <image.e01> --e01 -o timeline.csv
```

Build a DFIR timeline from a directory of VHDX images, useful when processing multiple VM snapshots.

```
$ hayabusa dfir-timeline -d <images_dir> --vhdx -o timeline.csv
```

Build a DFIR timeline from a single VHDX image without mounting it or needing admin rights

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a directory of mixed VHDX and E01 images in one pass to produce a combined DFIR timeline

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Collect log metrics from a directory of E01 images without mounting or extracting them

```
$ log-metrics --e01 -d images/ -U -o logs.csv
```

Build a DFIR timeline from a forensic E01 disk image without first extracting event logs manually.

```
$ hayabusa dfir-timeline -f <image.e01> --e01 -o timeline.csv
```

Build a DFIR timeline from a directory of VHDX images, useful when triaging multiple acquired VM disk snapshots.

```
$ hayabusa dfir-timeline -d <images_dir> --vhdx -o timeline.csv
```

Build a DFIR timeline from a VHDX disk image acquired from a compromised host, without mounting or extracting files.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a directory of mixed VHDX and E01 forensic images in one pass to generate a combined DFIR timeline.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Collect log metrics across a folder of E01 image segments without admin rights or FUSE.

```
$ log-metrics --e01 -d images/ -U -o logs.csv
```

Build a DFIR timeline from a single VHDX image without mounting it or needing admin rights.

```
$ hayabusa dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a directory of mixed VHDX and E01 images in one pass to produce a combined timeline.

```
$ hayabusa dfir-timeline --vhdx --e01 -d images/ -w -U -o
timeline.csv
```

Collect log metrics from a directory of E01 images to quickly assess log coverage before full analysis.

```
$ hayabusa log-metrics --e01 -d images/ -U -o logs.csv
```

Build a DFIR timeline directly from a forensic E01 disk image without first extracting event logs.

```
$ hayabusa dfir-timeline -f <image.e01> --e01 -o timeline.csv
```

Build a DFIR timeline from a VHDX disk image directory, useful when triaging a VM snapshot.

```
$ hayabusa dfir-timeline -d <DIR> --vhdx -o timeline.csv
```

Build a DFIR timeline from a single VHDX image without mounting it or needing admin rights.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Build a DFIR timeline from a directory of E01 images, letting Hayabusa find multi-segment sets automatically.

```
$ dfir-timeline --e01 -d images/ -w -U -p verbose -o
timeline.csv
```

Scan a mixed directory of VHDX and E01/Ex01 images in one pass to produce a combined forensic timeline.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Triage a forensic E01 image acquired from a compromised host without mounting it or extracting logs to disk.

```
$ hayabusa csv-timeline --e01 <image.e01> -o timeline.csv
```

Recursively find and scan all VHDX images in an evidence directory up to three subdirectory levels deep, preserving full evidence paths in the output.

```
$ hayabusa csv-timeline --vhdx <evidence-dir/> -o timeline.csv
```

Build a full DFIR timeline from a single VHDX image without mounting it or needing admin rights.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Process a mixed directory containing both VHDX and E01/Ex01 images in one pass to produce a combined timeline.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Build a full DFIR timeline from a VHDX image without mounting it or needing admin privileges

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a directory of mixed VHDX and E01 forensic images in one pass to generate a combined timeline

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Collect log metrics across a folder of E01 disk images without mounting or extraction

```
$ log-metrics --e01 -d images/ -U -o logs.csv
```

Build a DFIR timeline from a VHDX disk image without extracting event logs first.

```
$ hayabusa dfir-timeline -d /mnt/evidence --vhdx
```

Build a DFIR timeline from a single E01 forensic image file.

```
$ hayabusa dfir-timeline -f /cases/disk.e01 --e01
```

Build a DFIR timeline from a single VHDX image without mounting it or needing admin rights.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Build a DFIR timeline across a directory of mixed VHDX and E01 images in one pass.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Collect log metrics from a directory of E01 forensic images, e.g. to triage which hosts have relevant event logs before full analysis.

```
$ log-metrics --e01 -d images/ -U -o logs.csv
```

During a forensic investigation, scan a acquired EWF image directly for Windows event-log detections without extracting or mounting it.

```
$ hayabusa scan --e01 <path/to/image.E01>
```

Triage a VHDX disk image recursively across all contained NTFS volumes, preserving evidence paths in the output.

```
$ hayabusa scan --vhdx <path/to/image.vhdx>
```

Build a DFIR timeline from a single VHDX image on any platform without mounting it or needing admin privileges.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Build a DFIR timeline from a physical E01 disk image, pointing at the first segment and letting Hayabusa locate the rest automatically.

```
$ dfir-timeline --e01 -f disk.E01 -w -U -p verbose -o
timeline.csv
```

Process a directory that contains a mix of VHDX and E01 images in one pass to build a combined timeline.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Build a DFIR timeline directly from a forensic E01 disk image without first extracting event logs.

```
$ hayabusa dfir-timeline -f <image.e01> --e01 -o timeline.csv
```

Build a DFIR timeline from a VHDX disk image by pointing at a directory of images.

```
$ hayabusa dfir-timeline -d <images_dir> --vhdx -o timeline.csv
```

Build a DFIR timeline from a single VHDX image without mounting it or needing admin rights

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a mixed directory containing both VHDX and E01/Ex01 images and produce a combined timeline

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Build a DFIR timeline directly from a forensic disk image without mounting or extracting event logs manually.

```
$ hayabusa dfir-timeline -f <image.e01> --e01 -o timeline.csv
```

Build a DFIR timeline from a VHDX image on any OS without mounting it or needing admin rights.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Build a DFIR timeline from an E01 disk image, letting Hayabusa find the remaining segments automatically.

```
$ dfir-timeline --e01 -f disk.E01 -w -U -p verbose -o
timeline.csv
```

Scan a directory of mixed VHDX and E01 images in a single pass during bulk triage.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Build a full DFIR timeline from a single VHDX image without mounting it or requiring admin rights.

```
$ dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o
timeline.csv
```

Scan a directory of E01 images and produce log metrics — useful for rapid triage across multiple acquired drives.

```
$ log-metrics --e01 -d images/ -U -o logs.csv
```

Analyse a mixed directory containing both VHDX and E01/Ex01 images in one pass.

```
$ dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Triage a seized EWF forensic image directly on a Linux analyst workstation without extracting logs or mounting the image first.

```
$ hayabusa csv-timeline --e01 <image.E01> --output timeline.csv
```

Scan an entire directory of VHDX images up to three subdirectory levels deep — useful when processing a batch of VM disk images from an IR collection.

```
$ hayabusa csv-timeline --vhdx <images_dir/> --output
timeline.csv
```

Build a DFIR timeline from an E01 forensic image without first extracting event logs to disk.

```
$ hayabusa dfir-timeline -f /cases/evidence.e01 --e01 -o
timeline.csv
```

Run Hayabusa config against a forensic disk image instead of extracted .evtx files, useful when you have a raw acquisition.

```
$ hayabusa config -d <DIR> --vhdx
```

Target a single E01 forensic image directly with the config command.

```
$ hayabusa config -f <FILE> --e01
```

Build a full DFIR timeline from a VHDX image without mounting it or needing admin rights — useful during triage of a collected VM disk.

```
$ hayabusa dfir-timeline --vhdx -f host.vhdx -w -U -p verbose -o timeline.csv
```

Run log-metrics across a directory of E01 acquisition images to quickly assess log coverage across multiple evidence files.

```
$ hayabusa log-metrics --e01 -d images/ -U -o logs.csv
```

Process a mixed evidence directory containing both VHDX and E01 images in one timeline pass.

```
$ hayabusa dfir-timeline --vhdx --e01 -d images/ -w -U -o timeline.csv
```

Build a DFIR timeline directly from a forensic E01 image without extracting event logs first.

```
$ hayabusa dfir-timeline -f <image.e01> --e01
```

Build a DFIR timeline from a directory of VHDX disk images in a batch acquisition workflow.

```
$ hayabusa dfir-timeline -d <images_dir> --vhdx
```

— Flags, commands, path format and unsupported cases all named with runnable examples

product docs

02 **Shallow-clone rules fetching in update-rules** Not checked yet

IMPROVED 67

Running `update-rules` on a fresh install now performs a shallow clone of the Sigma rules repository, cutting the initial download from ~37 MB of full history down to ~8 MB; subsequent runs still fetch and fast-forward as before.

— Gives before/after size numbers but no new flag to invoke v4.5.0

THINNER COVERAGE BELOW

03 **Rules-free commands run without spurious errors** Not checked yet

IMPROVED 55

`update-rules`, `list-profiles`, `list-contributors`, `set-default-profile`, and `sort-csv` — commands that do not use detection rules — now run cleanly without a `rules` folder present, instead of printing spurious Cannot open file errors.

— Names five affected commands but only describes an error-message fix v4.5.0

---

04 **Version shown at startup** Not checked yet IMPROVED 23

Hayabusa now prints its version number above Start time when it starts.

— Minor cosmetic output change with no further detail v4.5.0

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ v2026.5.0 NOTES ↗ CODE ↗ RANK

iLEAPP is an open-source iOS forensic artifact parsing tool that extracts and analyzes data from iPhone and iPad device backups or images.

> iLEAPP v2026.5.0 adds decryption of Find My beacon tracking artifacts, new Google Keep/Tasks/Maps parsers, several social and browser artifacts, and a major consolidation of artifact categories alongside a more resilient iOS version detection fallback.

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

01 **Find My beacon decryption and tracking artifacts** Not checked yet

NEW 85

New `searchpartyd` artifacts now decrypt and surface Beacon Estimated Locations, Owned Beacons, Shared Beacons (from `ItemSharingKeys.db`), Unwanted Tracker Alerts and Locations, Observed Advertisements (from `Observations.db`, with recovered deleted rows), and Beacon Daily Addresses, using the `BeaconStore` keychain key for decryption.

— Names exact source databases and keychain key, but no CLI step given v2026.5.0

02 **iOS 27 shared albums artifact (Ph036)** Not checked yet NEW

70

New artifact `Ph036` extracts iOS 27 shared albums, posts, comments, and reactions from `Photos.sqlite`.

— Names artifact ID and source database clearly v2026.5.0

03 **Artifact categories consolidated into 31 groups** Not checked yet

IMPROVED 60

Artifact categories are folded from 262 into 31 purpose-based groups (Messaging, Browsers, etc.), while Biome, KnowledgeC, Unified Logs, PowerLog, Sysdiagnose, Photos.sqlite, and File System are kept whole.

— Gives before/after counts and named exceptions, no UI path v2026.5.0

04 **iOS version fallback from diagnostic reports** Not checked yet

IMPROVED 60

Extractions with no iOS version file now fall back to reading the version from the newest `.ips` diagnostic report, so version-gated artifacts still run.

— Explains mechanism and trigger condition clearly v2026.5.0

THINNER COVERAGE BELOW

05 **Google Keep, Tasks, and Maps artifacts** Not checked yet NEW

50

New artifacts extract Google Keep notes, list items, annotations, media, and tasks, plus Google Tasks from Keep, Gmail, Chat, and Calendar, and Google Maps Recent Searches.

— Lists data types parsed but no file or DB names given v2026.5.0

06 **TikTok, Snapchat, and Chromium Favicons artifacts** Not checked yet

NEW

45

New TikTok friends feed and follow feed played items artifacts are now parsed, alongside a new Snapchat Friending Badges artifact and a new Chromium Favicons artifact covering Chrome, Brave, and Edge.

— Named apps and artifact types but no extraction detail v2026.5.0

WAS THIS USEFUL? 👍 👎

ALEAPP

Feedback ⓘ v2026.5.0 NOTES ↗ CODE ↗

ALEAPP is an open-source Android log events and protobuf parser used by digital forensic investigators to extract and analyze artifacts from Android devices.

ALEAPP adds a six-step GUI wizard with dark/light themes, reads tool metadata from `data/leapp_config.json`, and collapses 380 artifact categories into 30.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.5.0
https://github.com/abrignoni/ALEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.5.0
```

New six-step GUI wizard (Case, Source, Artifacts, Review, Run, Results) with dark and light themes; command-line usage is unchanged, and the wizard can stop a run after the current artifact finishes. Not checked yet

Read ALEAPP's name, version, and description from `data/leapp_config.json` instead of hard-coded values. Not checked yet

The 380 artifact categories are folded into 30 grouped categories in the artifact picker. Not checked yet

Eight artifacts now read every differing copy of a database across `data/data`, `data/user/0`, and `data_mirror`. Not checked yet

NEW DETECTION WhatsApp Business artifacts: WhatsApp Business contacts

NEW DETECTION WhatsApp and WhatsApp Business media file artifacts: WhatsApp Business media files

NEW DETECTION Passpoint Providers artifact

NEW DETECTION Content query text artifacts from android_triage and ALEX: ALEX raw text output

UPDATED DETECTION Digital Wellbeing multi-user store coverage: app_usage

UPDATED DETECTION Google Tasks extended source coverage: Gmail, Calendar, Keep, Google Chat

UPDATED DETECTION Notification History icon type and image reporting

UPDATED DETECTION Private Photo Vault extended fields

UPDATED DETECTION what3words inactive Realm version row reporting

RETIRED DETECTION Android WeChat artifact removed

└──▷ **BREAKING ON UPGRADE**

! The `aleappGUI` executable is removed; `aleapp` (without arguments) now opens the window instead.

! The Android WeChat artifact is removed and will produce no output on upgrade.

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

YARA-X

Feedback ⓘ v1.25.0 NOTES ↗ RANK

YARA-X is a rewrite of YARA in Rust, used to create and apply pattern-matching rules for identifying and classifying malware samples.

> YARA-X v1.25.0 adds missing constants to the `elf` and `pe` modules and breaks backward compatibility with rules serialized under the previous `bincode` format.

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

01 New constants in elf and pe modules

Not checked yet NEW

63

The `elf` module gains ELF section flag constants `SHF_WRITE`, `SHF_ALLOC`, `SHF_EXECINSTR`, and others for use in rules. The `pe` module gains previously missing PE machine type constants `MACHINE_TARGET_HOST`, `MACHINE_SH3E`, `MACHINE_TRICORE`, `MACHINE_CEF`, and `MACHINE_CEE`.

– Names all constants but no usage example shown v1.25.0

THINNER COVERAGE BELOW

02 Bincode serialization format incompatibility

Not checked yet

45

BREAKING

Rules compiled with the previous `bincode -based` serialization format can no longer be loaded; all previously serialized rules must be recompiled.

– Clear migration note but no detail on cause v1.25.0

└──▷ BREAKING ON UPGRADE

! Rules compiled with the previous `bincode -based` serialization format can no longer be loaded — all previously serialized rules must be recompiled.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Docker Compose

Feedback ⓘ v5.8.1 NOTES ↗ RANK

Docker Compose is a tool for defining and running multi-container Docker applications using a single YAML configuration file.

> Docker Compose v5.8.1 speeds up builds by loading OCI images directly from the local cache instead of re-pulling them.

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

01 **Direct OCI cache loading for images** Not checked yet IMPROVED

33

Docker Compose now loads OCI images directly from the local cache when possible, skipping redundant pulls.

— Describes behaviour but no flag, command or metric given v5.8.1

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.53.0 NOTES ↗ CODE ↗ RANK

Prowler is an open-source security tool that assesses AWS, Azure, and GCP cloud environments against security best practices and compliance frameworks.

> Prowler 5.53.0 adds support for custom and sovereign Azure/M365 cloud regions, a new check for overly permissive GitHub Actions OIDC trust in AWS IAM roles, Posture Change PDF reporting, and a new AI Security posture page in Prowler Cloud.

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

01 Custom and sovereign Azure/M365 cloud regions

Not checked yet NEW

90

Scan custom, sovereign, or air-gapped Azure clouds with the new `AzureCustom` region using `--azure-authority-host`, `--azure-resource-manager-endpoint`, `--azure-graph-endpoint`, and `--azure-logs-endpoint` CLI flags. Prowler Private Cloud configures Azure and M365 provider regions with `PROWLER_AZURE_REGION` (`AzureCloud`, `AzureUSGovernment`, `AzureChinaCloud`, or `AzureCustom`) and `PROWLER_M365_REGION` (`M365Global`, `M365USGovernment`, `M365China`, or `M365Custom`), with custom endpoints set via `PROWLER_AZURE_AUTHORITY_HOST`, `PROWLER_AZURE_RESOURCE_MANAGER_ENDPOINT`, `PROWLER_AZURE_GRAPH_ENDPOINT`, and `PROWLER_AZURE_LOGS_ENDPOINT`.

— Names every flag and env var for both CLI and Private Cloud. 5.53.0

02 GitHub Actions OIDC trust check for IAM roles

Not checked yet NEW

90

New `iam_role_github_oidc_trust_subject_restricted` check detects overly permissive GitHub Actions OIDC trust in AWS IAM roles, and the `iam_github_oidc_allowed_owners` configuration option limits accepted GitHub organisation owners.

Restrict which GitHub organisation owners are accepted when scanning for permissive OIDC trust in IAM roles.

```yaml
iam_github_oidc_allowed_owners:
  - my-org
  - trusted-partner-org
```

— Names check id, config key, plus a usage example. 5.53.0

03 **Posture Change PDF reports** Not checked yet NEW 85

Generate a Posture Change PDF report, in English or Spanish, showing how a provider's security posture changed between two dates, via `POST /api/v1/posture-change-reports` or the Custom Scans and Reports page.

— Names exact endpoint and UI path. 5.53.0

04 **AI Security posture page in Prowler Cloud** Not checked yet NEW 75

New AI Security page covers model vendors (Claude, OpenAI, ChatGPT) and cloud AI services (Amazon Bedrock, Azure AI, Vertex AI) with posture summaries and failed findings by severity, against six AI compliance frameworks: OWASP Top 10 for LLM Applications 2025, OWASP Top 10 for Agentic Applications 2026, OWASP MCP Top 10 2025, NIST AI RMF 1.0, ISO/IEC 42001:2023, and the EU AI Act.

— Names all six frameworks and services, but only a UI path. 5.53.0

05 **Resource evidence appendix in compliance PDFs** Not checked yet IMPROVED 65

Compliance PDF reports (ENS, ThreatScore, NIS2, CSA CCM, CIS) now include a Resource evidence appendix with the full configuration of affected resources, matching the Evidence tab of a finding.

— Names all five compliance frameworks affected, UI only. 5.53.0

06 **Premium upgrade prompts in Prowler Local Server** Not checked yet NEW 60

Prowler Local Server now surfaces a Premium sidebar section, an upgrade badge, and a comparison page showing which capabilities require Prowler Cloud or Private Cloud, with preview pages for Exploitability issues and Realtime instead of 404s.

— UI-only change, no config or API named. 5.53.0

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - October 20, 2026 BLOG ↗

RANK

Sysdig Secure is a cloud-native security platform that detects threats, manages vulnerabilities, and ensures compliance across containers, Kubernetes, and cloud environments.

> Sysdig Secure introduced a Technical Preview Supply Chain dashboard that summarises license and end-of-life findings across pipeline, registry, and runtime resources.

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

01 **Supply Chain dashboard (Technical Preview)** Not checked yet NEW

65

A new dashboard summarises license and end-of-life findings across pipeline, registry, and runtime resources. Users can navigate to `Dashboards > Supply Chain`, apply Zone and Stage filters, and select a panel or image tile to drill into the matching findings on the Supply Chain Findings page.

Review end-of-life OS exposure across your supply chain — navigate to the Supply Chain dashboard, filter by zone and stage, then select any image tile to drill into its findings.

> 📍 In Sysdig Secure, go to Dashboards › Supply Chain, apply the Zone and Stage filters, then select a panel or image to open the matching findings on the Supply Chain Findings page.

— Clear UI path given, but underlying mechanism unexplained

Secure SaaS - November 26, 2026

NEW SUPPORTED PLATFORM arm64 architecture support for Agentless Workload Scanning: AWS ECS Fargate, Lambda Functions, EKS, GCP GKE, Azure AKS

> Secure SaaS - October 20, 2026 BLOG ↗ LATEST STABLE

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Azure network checks for CIS v7.0.0: NOTES ↗

RANK

Plerion is a cloud security platform that provides risk-based prioritization and remediation guidance to help organizations manage and reduce their cloud attack surface.

> Plerion's latest release adds 14 new Azure network checks aligned to CIS Microsoft Azure Foundations Benchmark v7.0.0, expands asset collection to five new Azure resource types, and refines several existing Azure and AWS checks, with some severity downgrades and a check retirement along the way.

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

01 **PLERION-AZURE-294 now inspects port 443, retires -237**

Not checked yet 70

BREAKING

Check `PLERION-AZURE-294` now inspects port 443 in addition to port 80, so most groups serving the public will fail. In code scans, the retired check `PLERION-AZURE-237` is merged into `PLERION-AZURE-294` and its findings become `PLERION-AZURE-294` findings; exemptions set on `PLERION-AZURE-237` do not carry over.

— Names check IDs, scope change and exemption migration impact

Azure network checks for CIS v7.0.0:

02 **14 new Azure network checks for CIS v7.0.0** Not checked yet NEW

60

Plerion adds 14 new Azure network checks (`PLERION-AZURE-555` through `PLERION-AZURE-568`) covering network security groups, subnets, DDoS protection, Bastion, Application Gateway, Front Door, and VPN gateways, aligned to the CIS Microsoft Azure Foundations Benchmark v7.0.0. Four of these checks also run on code scans, extending infrastructure-as-code coverage to Azure network controls.

— Named check range and CIS alignment, no usage path given

Azure network checks for CIS v7.0.0:

03 **API Gateway WAF exposure check improvements** Not checked yet

IMPROVED 60

Check `PLERION-AWS-564` now reads the web ACL on each stage of a REST API, passing a stage only when the ACL restricts access to a small portion of the internet (such as company IP ranges), and reads IP sets from the AWS Firewall Manager admin account when Firewall Manager is in use.

— Detailed mechanism including Firewall Manager integration

Azure network checks for CIS v7.0.0:

THINNER COVERAGE BELOW

04 **Live-subscription scanning for five Azure checks** Not checked yet

IMPROVED

55

Checks `PLERION-AZURE-197`, `PLERION-AZURE-199`, `PLERION-AZURE-200`, `PLERION-AZURE-201`, and `PLERION-AZURE-202`, which previously ran only on code scans, now also run against live Azure subscriptions.

— Named checks with clear before/after scope change Azure network checks for CIS v7.0.0:

05 **Severity lowered on several Azure checks** Not checked yet

BREAKING

55

The severity of `PLERION-AZURE-85`, `PLERION-AZURE-86`, `PLERION-AZURE-88`, `PLERION-AZURE-90`, `PLERION-AZURE-294`, and `PLERION-AZURE-300` is lowered from High to Low; `PLERION-AZURE-89`, `PLERION-AZURE-201`, and `PLERION-AZURE-202` are lowered from Medium to Low. Custom severities set by users are preserved.

— Names all affected checks and severity before/after

Azure network checks for CIS v7.0.0:

06 **Per-Network-Watcher flow log check** Not checked yet IMPROVED

50

Check `PLERION-AZURE-300` now produces a finding on each Network Watcher individually and fails when a flow log does not send to a Log Analytics workspace, replacing the previous manual check.

— Describes mechanism change for a single named check

Azure network checks for CIS v7.0.0:

07 **Five new Azure asset types collected** Not checked yet NEW

40

Plerion now collects five new Azure asset types: Application Gateways, Application Gateway web application firewall policies, Front Door profiles, Front Door firewall policies, and virtual network gateways.

— Lists new asset types with no further mechanism Azure network checks for CIS v7.0.0:

└──▷ BREAKING ON UPGRADE

! In code scans, check `PLERION-AZURE-237` is retired and its findings become `PLERION-AZURE-294` findings; exemptions set on `PLERION-AZURE-237` do not carry over.

! The severity of `PLERION-AZURE-85`, `PLERION-AZURE-86`, `PLERION-AZURE-88`, `PLERION-AZURE-90`, `PLERION-AZURE-294`, and `PLERION-AZURE-300` is lowered from High to Low;

`PLERION-AZURE-89`, `PLERION-AZURE-201`, and `PLERION-AZURE-202` are lowered from Medium to Low. Custom severities are preserved.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

Filter Deployment/continuous events by violation reason (2026-11-29) BLOG ↗

RANK

Trend Vision One is a unified cybersecurity platform that integrates threat detection, response, and risk management across endpoints, networks, and cloud environments.

> Trend Vision One shipped two targeted usability improvements this window: a violation-reason filter for Container Security Findings and customizable status aliases in Case Management.

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

01 Violation reason filter in Container Security Findings

Not checked yet NEW 72

Deployment/continuous events (Kubernetes and Amazon ECS) in Container Security Findings can now be filtered by one or more violated policy rules, using Match any or Match all logic, via Cloud Security → Container Security → Findings.

Narrow down which Kubernetes or ECS deployment events violated a specific set of policy rules — useful when triaging a noisy findings list after a policy change.

> 📍 In the console, go to Cloud Security → Container Security → Findings, open the Violation reason filter, select one or more policy rules, then choose Match any to see events that violated at least one rule or Match all to see events that violated every selected rule.

— Names exact UI path and filter logic, no mechanism beyond that

Filter Deployment/continuous events by …

02 Custom status aliases in Case Management Not checked yet NEW

64

Case Management now supports assigning custom status aliases of up to 30 characters to the To do, In progress, and Closed case statuses, configured under Workflow and Automation → Case Management.

Rename default case statuses to match your SOC's terminology so analysts see familiar terms without changing underlying workflow logic.

> 📍 In the console, go to Workflow and Automation → Case Management, open the status settings, and assign an alias (up to 30 characters) to To do, In progress, or

> Closed.

— Exact path, character limit, and status names given

Status aliases now available in Case Ma…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Sekoia.io

Feedback ⓘ changelog · Nov 28, 2026 CHANGELOG ↗ RANK

Sekoia.io is a cloud-native cybersecurity platform offering threat intelligence, SIEM, and SOAR capabilities to detect, investigate, and respond to cyber threats.

> Sekoia.io's latest changelog adds automation controls and analytics to Elevate, risk scoring and vulnerability discovery to Reveal, new SOL data sources for monitoring playbook health, bulk case management in Defend, curated threat briefings in Intelligence, and a reworked Playbooks module.

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

01 Activation filters for Elevate automation

Not checked yet NEW

80

Sekoia Elevate adds activation filters for rule effort level (Elementary, Intermediate, Advanced) so you can choose which alerts Elevate analyzes automatically, keeping automated runs targeted at high-signal alerts rather than ones playbooks already close. Set the filter in Elevate Activation Settings.

Scope Elevate to only analyze high-signal alerts by setting the rule effort level filter in Elevate Activation Settings, keeping automated runs away from noise that playbooks already handle.

> 📍 In Sekoia Elevate, go to Activation Settings, set the Rule effort level filter to the desired levels (Elementary, Intermediate, or Advanced), and save.

— Names exact filter levels and UI setting, with step-by-step example

changelog-20261128-29e19170

02 SOL playbook data sources for automation health

Not checked yet

NEW

75

Three new SOL data sources — `playbooks`, their runs, and every step they execute — are now queryable in dashboards to monitor automation health. You can build a query filtered by run status or step outcome and pin the result to a dashboard.

Track automation health by querying the new `playbooks` SOL data source in a dashboard to see run status across all your playbooks at a glance.

> 📍 In Sekoia Defend, open SOL, select the `playbooks` data source, build your query to filter by run status or step outcome, then pin the result to a dashboard.

— Names exact data sources with a runnable query example changelog-20261128-29e19170

03 **Bulk case management in Sekoia Defend** Not checked yet IMPROVED

70

Cases in Sekoia Defend now support bulk actions to change priority, status, verdict, or assignment across many cases at once, plus case templates and custom status migration. The Cases list also adds a `Me` option in the `Assigned to` filter that stays relative per analyst in saved views, and a Type filter to show only AI cases or only manual ones.

— Names several concrete fields and filters but no command or path

changelog-20261128-29e19170

THINNER COVERAGE BELOW

04 **Analytics page in Sekoia Elevate** Not checked yet NEW 55

Sekoia Elevate adds an Analytics page showing time to verdict, false positives qualified, true positives qualified, and analyst agreement rates over the last 30 days.

— Names specific metrics but gives no navigation detail changelog-20261128-29e19170

05 **Revamped Playbooks module in Defend** Not checked yet IMPROVED

50

The Playbooks module in Sekoia Defend now has a table view, guided creation, and version history, replacing the previous card-based listing.

— Lists UI changes but no path or configuration detail changelog-20261128-29e19170

06 **Curated threat briefings in Sekoia Intelligence** Not checked yet

IMPROVED 45

Sekoia Intelligence's country and sector pages now open on a curated threat briefing from Sekoia TDR, showing who is targeting a country or industry and with which malware, techniques, and vulnerabilities.

— Describes content shown but not how to configure or query it

changelog-20261128-29e19170

07 **Asset Risk Scoring and discovery in Reveal** Not checked yet NEW

40

Sekoia Reveal adds Asset Risk Scoring, giving every asset a risk score, and Application and Vulnerability Discovery, which lists applications installed on each host and flags vulnerable ones.

— Describes what each does but no mechanism or access path changelog-20261128-29e19170

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Value Overview for Mobot, SOC Analyst Agent, and Root Cause Agent (Manage) (2026-11-27

BLOG ↗ RANK

Sumo Logic is a cloud-based platform that aggregates, analyzes, and visualizes machine data and logs for security monitoring and operational insights.

> Sumo Logic introduces a Value Overview tab that quantifies time saved by its AI agents, giving analysts a dashboard to track, filter, and tune the assumptions behind those savings estimates.

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

01 **Value Overview tab for AI agent time savings** Not checked yet NEW

63

A new Value Overview tab shows combined and per-agent estimated hours saved across Mobot, SOC Analyst Agent, and Root Cause Agent, with a breakdown by agent. An agent metrics view lets analysts review time saved, investigation activity, and time to verdict per agent, while filters and date-range search narrow results by source and time window. Configurable assumptions let analysts edit the manual-time values used in the hours-saved calculations.

Review how much investigation time Mobot and the SOC Analyst Agent have saved your team over the past month, then adjust the assumed manual-review time per alert to reflect your team's actual baseline.

> 📍 In the Sumo Logic console, go to Value Overview › Combined Time Reclaimed, set a date range, filter by source, then select Configurable assumptions and edit the manual-time values to match your team's workflow.

— Names the tab, agents, and metrics but no API or config key

Value Overview for Mobot, SOC Analyst A…

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Google Gemini 4 Neon is now available on Unity Gateway (2026-11-25) BLOG ↗

RANK

Databricks is a unified data analytics platform that enables collaborative data engineering, machine learning, and analytics using Apache Spark in the cloud.

> Databricks expanded account-level access to Genie One for SCIM-provisioned users and added Google's Gemini Nano Banana 2.1 image generation model to Unity Gateway.

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

01 **Account-only user access to Genie One** Not checked yet NEW

60

Users provisioned through automatic identity management or account-level SCIM can now access Genie One without being assigned to a workspace or granted a workspace entitlement; Unity Catalog privileges govern data access.

— Explains mechanism and governance but no config steps shown

Account-only user access to Genie One i…

THINNER COVERAGE BELOW

02 **Gemini Nano Banana 2.1 on Unity Gateway** Not checked yet

NEW INTEGRATION NEW

50

Google Gemini Nano Banana 2.1 is now available on Unity Gateway, enabling image generation from text and image inputs via Foundation Model APIs.

— Names model and API surface but no usage example

Google Gemini Nano Banana 2.1 is now av…

> Google Gemini 4 Neon is now available on Unity Gateway (2026-11-25) BLOG ↗
>
> LATEST STABLE

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261127.0 NOTES ↗ RANK

OpenCTI is an open-source threat intelligence platform that helps organizations manage, store, and analyze cyber threat intelligence data.

> OpenCTI tightens startup security with default-credential checks and stricter 2FA enforcement, improves visibility into retention-rule and connector failures, and expands custom fields and shared-object UI with new search, decay, and enrichment options.

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

01 **Retention rule error logging and auditing** Not checked yet IMPROVED

52

Retention rule runs now log, audit, and expose their last error so operators can see why a rule failed without digging through raw logs.

– Explains behavior but no UI location or API field named 7.261127.0

02 **Default-credential checks for Elastic and Redis** Not checked yet

NEW

47

OpenCTI now checks whether Elastic and Redis are still using default credentials at startup, alerting operators before those services are reachable.

– Names Elastic and Redis but no alert mechanism or UI path given 7.261127.0

03 **Basic auth blocked for 2FA-enabled accounts** Not checked yet

47

BREAKING

API Basic authentication is now refused for accounts that have two-factor authentication enabled, enforcing stronger auth paths for those users.

– Clear behavior change but no error code or migration note given 7.261127.0

04 **Scheduled connector exception reporting** Not checked yet NEW

45

The Python client now reports exceptions raised during scheduled connector runs back to the platform, surfacing connector failures in the UI.

– Describes mechanism but not where in UI failures appear 7.261127.0

05 **Shared object enrichment and navigation** Not checked yet NEW

43

OpenXTM and OpenCRQ counterparts, pivot views, and `connect` calls to action are now available on every shared object.

— Names three UI additions but no screen path or workflow detail `7.261127.0`

---

06 **Custom fields gain search, decimal, and decay support**

`Not checked yet` **NEW** 42

Custom fields now support `full-text search`, `decimal numbers`, and `decay rules`.

— Names three additions but no config or usage detail `7.261127.0`

---

07 **Flag to reduce backend error forwarding** `Not checked yet` **NEW**

18

A flag is now available to reduce error forwarding from the backend.

— Flag mentioned but not named or located `7.261127.0`

WAS THIS USEFUL? 👍 👎

Lookyloo

Feedback ⓘ v1.43.2 NOTES ↗ RANK

Lookyloo is a web security tool that captures and visualises the full tree of requests made when loading a webpage, aiding malware and phishing analysis.

> Lookyloo's v1.43.2 release adds a forced capture option for authenticated users and redesigns the tree view capture page with new navigation tools.

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

01 **Redesigned tree view capture page** Not checked yet IMPROVED

55

The tree view capture page was redesigned with updated menus, a legend filter, and zoom and navigation controls.

— Names specific UI additions but no exact navigation path v1.43.2

02 **Forced capture for authenticated users** Not checked yet NEW

30

Authenticated users can now force a capture, per the release summary, though no further mechanism is described.

— Named but not explained beyond the summary line v1.43.2

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ changelog · Nov 28, 2026

CHANGELOG ↗ RANK

Horizon3 NodeZero is a horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

> Horizon3 NodeZero expanded Kubernetes Pentest configurability, added REST API coverage and parallelized test flows to WebApp pentest, and introduced an Audit Log dashboard, a Tripwires webhook connector, and a Client Management view for multi-tenant parent accounts.

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

01 **Kubernetes Pentest configuration options** Not checked yet NEW

75

Kubernetes Pentest adds options to configure a Custom Service Account, Import Rules from YAML to load custom rules from a file, and set up a proxy before launching a test.

Scope a Kubernetes pentest to use a specific service account and load custom rules from a YAML file rather than relying on defaults.

📍 In the NodeZero portal, open a Kubernetes Pentest, go to Options, set Custom Service Account, choose Import Rules from YAML, upload your rules file, and configure a proxy if needed before launching.

— Names exact UI options with a step-by-step usage example changelog-20261128-fbc6c9fa

THINNER COVERAGE BELOW

02 **Client Management view for parent accounts** Not checked yet NEW

55

Parent accounts in multi-tenant environments get a new top-level Client Management view showing clients' operational health, license allocations, and API keys.

— Names three data points shown but no navigation steps changelog-20261128-fbc6c9fa

03 **Audit Log dashboard for Org Admins** Not checked yet NEW

50

Org Admins get an Audit Log dashboard to review, investigate, and export records of account activity.

— Describes dashboard function, gives clear UI starting point changelog-20261128-fbc6c9fa

04 **Tripwires webhook connector** Not checked yet NEW 45

A new webhook connector lets Tripwires alerts be sent to SIEM, SOAR, or automation platforms.

— Names destination categories but no config detail changelog-20261128-fbc6c9fa

05 **REST API testing in WebApp pentest** Not checked yet NEW 45

WebApp pentest coverage expands to test REST APIs that do not expose a web UI.

— States the capability but no mechanism or scope changelog-20261128-fbc6c9fa

06 **Parallelized WebApp pentest execution** Not checked yet IMPROVED 45

WebApp pentest now runs parallelized test flows for faster results and produces reports with verifiable evidence and fewer false positives.

— Describes performance and quality improvement, no concrete numbers changelog-20261128-fbc6c9fa

07 **New vulnerability detection coverage** Not checked yet NEW 35

NodeZero adds detection coverage for vulnerabilities in LibreNMS, Drupal Commerce, IBM App Connect Enterprise, and IBM Langflow.

— Names four affected products but no CVE or mechanism detail changelog-20261128-fbc6c9fa

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Attack Surface Management

Axonius

Feedback ⓘ Axonius Release Notes 10.0.3 (2026-11-22)

BLOG ↗

Axonius is a cybersecurity asset management platform that discovers, inventories, and manages all devices, users, and cloud assets across an organization's environment.

Axonius 10.0.3 adds credit-free AI Companion access, a Percentage function for Field Summary charts, Connected Apps and Granted Permissions in SaaS profiles, and merges Managed Identities into the Users asset.

The AI Companion no longer requires Axonius AI Credits and is now governed by a fair-use model, making it available to all customers without credit assignment, usage tracking, or capacity planning. Not checked yet

Scheduled AI Worker and Virtual Employee runs now read the latest daily asset snapshot instead of live asset data, with answers up to approximately 24 hours old. Not checked yet

The Manage Vulnerability Exception and Set Security Findings Rule Status agent tools now always require manual human approval and can no longer be set to auto-approve. Not checked yet

The AI Agent can now find enforcement sets and workflows that use or update a specific field, referenced by its title (custom, device, or user field). Not checked yet

The AI Agents chatbot can now explain Attack Path Analysis options, including scenarios, graph controls, node details, settings, and can read the scenario graph of a Security Finding. Not checked yet

+15 more

NEW INTEGRATION 11 new adapters

UPDATED INTEGRATION 72 updated adapters

UPDATED INTEGRATION Adapter trust tier changes: Wiz, Tenable Vulnerability Management, Qualys Cloud Platform, Qualys Container Security, Rapid7 InsightVM and 1 more

└──▷ BREAKING ON UPGRADE

! The `managed_users` asset type is deprecated; API v2 responses include Deprecation and Sunset headers with a sunset date of December 13, 2026 — switch to Users with the managed filter before that date.

! After upgrading, any existing Manage Vulnerability Exception or Set Security Findings Rule Status agent tool configurations set to auto-approve will default to requiring human approval.

! If the 'Hide security findings that were fixed more than X days ago' setting under

Status Resolver was empty (no limit), it now defaults to 14 days after upgrade.

- ! The CUs Behind Latest field is renamed to B-Releases Behind (CUs) and Preferred CUs Behind Latest is renamed to Preferred B-Releases Behind (CUs).
- ! Data scopes that include SaaS Applications now show only the SaaS child assets of the SaaS applications in the scope — review existing data scopes.
- ! Managed Identities no longer appear in navigation, the Assets side panel, global search results, or as an option for new charts, Enforcement Sets, the Queries module filter, and Asset Graph tiles.
- ! Agent Coverage counts in the Uncategorized Devices category and other categories may go down after upgrade.

WAS THIS USEFUL? 👍 👎

◆ **Vulnerability & Exposure Management**

Edgescan

Feedback ⓘ Release Notes - 3rd Nov 2026 NOTES ↗ RANK

Edgescan is a cloud-based vulnerability management platform that continuously assesses and prioritises security risks across web applications, networks, and APIs.

> Edgescan streamlined authenticated scan setup by letting users create a browser recording directly inside the Authentication Workflow wizard.

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

01 **In-wizard browser recording creation** Not checked yet NEW

55

The Authentication Workflow wizard's Browser Recording picker now supports creating and applying a new browser recording directly inside the wizard, removing the need to leave the wizard to upload one separately.

Create a browser recording mid-configuration when setting up authenticated scanning, without navigating away from the wizard.

📍 In the Authentication Workflow wizard, open the Browser Recording picker and select Create new recording to record and apply a login flow in one step.

— Clear UI path given, but limited mechanism detail. Release Notes - 3rd Nov 2026

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Clerk

Feedback ⓘ changelog · Nov 28, 2026 CHANGELOG ↗ RANK

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

> Clerk shipped metadata-based user search across API, SDK, and Dashboard surfaces, a native app CLI bootstrapper for iOS/Android/Expo, Native to Web SSO, and a breaking 1.0 release of its MCP authorization SDK.

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

01 **Native app CLI setup** Not checked yet NEW 95

The Clerk CLI adds `clerk init --starter` with `--framework ios`, `--framework macos`, or `--framework android` (plus `--prebuilt-auth-ui` to include the sign-in screen) to create native iOS, Android, or Expo apps with Clerk pre-wired; `--sign-in-with-apple` and `--sign-in-with-google` to enable native Sign in with Apple or native Sign in with Google on Android; `clerk doctor` to check iOS, macOS, or Android Clerk setup at any time; `--mode agent --json` to run agents through CLI setup non-interactively; and running `clerk init` inside an Xcode or Android project to link your Clerk application, register the app, and add the Clerk SDK automatically.

Bootstrap a new iOS app with Clerk authentication and the prebuilt sign-in screen in one command, then follow the iOS quickstart.

```
$ npx -y clerk@latest init --starter --framework ios --prebuilt-auth-ui
```

— Every flag named verbatim with a runnable command changelog-20261128-4cd2e60d

02 **MCP authorization SDK reaches 1.0 with breaking changes**

Not checked yet 95

BREAKING

`@clerk/mcp-tools` 1.0.0 adds the `publicAccess` name option to control MCP server discoverability per tool (or `discoverable: false` to disable it entirely), publishes an MCP Server Card and AI Catalog so clients and registries can discover your server, and includes a built-in `submit_feedback` tool to prompt agents to report problems and continue tasks. The release is breaking: MCP servers are now discoverable by default, letting clients without an access token connect and list tools, resources, resource templates, and prompts, unless `discoverable: false` restores the 0.9 behaviour; exports drop the Clerk suffix with the old names removed and no aliases; Clerk-verified tokens must carry your `resourceUrl` as their audience with stricter checks for custom

token verifiers; minimum Clerk SDK and MCP SDK versions are raised; and sessions stored by MCP clients before 1.0.0 must reconnect once after upgrading.

— Names every config option and migration consequence of the breaking change

changelog-20261128-4cd2e60d

03 **Paginated user listing without total count** Not checked yet NEW

60

getUserListWithoutCount() (JavaScript) and `user.ListWithoutCount` (Go), available in `@clerk/backend` v3.32.0 or `clerk-sdk-go` v2.10.0, skip computing the per-page user total to speed up paginated reads.

— Names the functions but gives no usage example or numbers changelog-20261128-4cd2e60d

04 **Native to Web SSO for mobile apps** Not checked yet NEW

60

Clerk adds Native to Web SSO for iOS, Android, and Expo apps, opening web pages with the user already signed in via a one-time link that expires after 60 seconds.

— Explains mechanism and expiry but no API/flag to invoke it changelog-20261128-4cd2e60d

└──▷ **BREAKING ON UPGRADE**

! MCP servers built with `@clerk/mcp-tools` 1.0.0 are discoverable by default: clients without an access token can connect and list tools, resources, resource templates, and prompts. Set `discoverable: false` to restore the 0.9 behaviour of requiring a token for every request.

! Exports from `@clerk/mcp-tools` drop the `Clerk` suffix and the old names are removed with no aliases.

! Clerk-verified tokens must carry your `resourceUrl` as their audience, and custom token verifiers receive stricter checks.

! The minimum Clerk SDK and MCP SDK versions are raised in `@clerk/mcp-tools` 1.0.0.

! Sessions stored by MCP clients before `@clerk/mcp-tools` 1.0.0 must reconnect once after upgrading.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-27) BLOG ↗

RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud's latest release lets admins run Workflows through natural-language MCP commands, override SCIM app display names without renaming groups, retrieve Mac device-lock PINs from the Admin Portal, and upgrade to a new Directory Insights API Python SDK with multi-org support.

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

01 Directory Insights API Python SDK 1.1.0 multi-org context

Not checked yet IMPROVED 92

The Python SDK for the Directory Insights API is updated to version 1.1.0, now passing organization context via the `x_org_id` argument on each method call instead of setting it on the API key configuration. Multi-org context can be passed as a comma-separated list, e.g. `idA,idB,idC`, in the `x-org-id` header, as shown by insights_client.get_events(x_org_id='idA,idB,idC', ...).

Pass multi-org context to the Directory Insights API Python SDK 1.1.0 to query events across several managed organizations in one call.

```python
insights_client.get_events(x_org_id='idA,idB,idC', ...)
```

— Names exact SDK version, argument, header, and runnable code

Release Notes 2026 (2026-11-27)

02 Per-app display name override for SCIM groups Not checked yet NEW

80

SCIM group provisioning adds a per-application Show In App As field, letting admins override the display name sent to a SCIM-connected app while keeping the internal JumpCloud group name unchanged. The override is set on the application's User Groups tab, and clicking Save and sync pushes the new name to the app immediately.

Override the display name sent to a SCIM-connected app without renaming the JumpCloud group, then push the change immediately.

> 📍 On the application's User Groups tab, enter the desired name in Show In App As, then click Save and sync to update the application right away.

— Names exact field and UI path but no API-level detail Release Notes 2026 (2026-11-27)

THINNER COVERAGE BELOW

03 **Workflows management via MCP Server for Admins** Not checked yet

NEW

53

JumpCloud's MCP Server for Admins now supports creating, modifying, duplicating, triggering, and troubleshooting Workflows via a connected MCP-enabled AI client using natural language, with no additional connection or credentials required.

— Describes capability but no exact commands or endpoints given

Release Notes 2026 (2026-11-27)

04 **Escrow PIN retrieval for locked Mac devices** Not checked yet NEW

53

After issuing a remote Device Lock command on a Mac, admins can retrieve the auto-generated escrow PIN from the Admin Portal by navigating to Highlights > Device Lock > View PIN.

Retrieve a locked Mac's escrow PIN from the Admin Portal after issuing a remote Device Lock command.

> 📍 In the Admin Portal, go to Highlights › Device Lock › View PIN to retrieve the auto-generated lock PIN for a locked Mac.

— Only described via example steps, no broader mechanism given

Release Notes 2026 (2026-11-27)

WAS THIS USEFUL? 👍 👎

Keycloak

Feedback ⓘ

Keycloak Workflows in Practice, Part 1: Automating Identity Lifecycle (2026-11-27)

BLOG ↗ RANK

Keycloak - the open source identity and access management solution. Add single-sign-on and authentication to applications and secure services with minimum effort.

> Keycloak's big addition this window is a new event-driven Workflow engine for automating identity lifecycle tasks, defined in YAML or JSON and managed via a dedicated Admin REST API and Admin Console UI.

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

01 **Workflow engine for identity lifecycle automation** Not checked yet

NEW 100

Keycloak adds an event-driven Workflow engine, defined in YAML or JSON, that triggers on realm events such as `user-created`, `user-authenticated`, user-role-granted(admin), `client-created`, and `client-authenticated`. Workflows are created via `POST /admin/realms/{realm}/workflows`, accepting `Content-Type: application/yaml` or `Content-Type: application/json`. Execution can be scoped with `if` conditions like has-user-attribute(department=engineering) or not has-user-attribute(permanent-admin=true), individual steps delayed with the `after` key (e.g. `after: 90d`, `after: 1d`, `after: 30s`), and overlapping runs controlled via `concurrency: restart-in-progress: true`. Built-in user-workflow steps include `notify-user`, `grant-role`, `revoke-role`, `join-group`, `leave-group`, `set-user-attribute`, `remove-user-attribute`, `add-required-action`, `remove-required-action`, `disable-user`, `logout-user`, `delete-user`, `unlink-user`, `invite-user`, and `restart`, plus client-workflow steps `delete-client` and `disable-client`. The background step-runner's polling frequency for due scheduled steps can be tuned at startup with `--spi-events-listener--workflow-event-listener--step-runner-task-interval`.

Automatically grant a role and add a user to a group the moment they are created with a matching department attribute — no manual admin steps needed.

```
$ curl -X POST
"http://localhost:8080/admin/realms/acme/workflows" \
  -H "Authorization: Bearer $TOKEN" \
  -H "Content-Type: application/yaml" \
  -d '
name: Onboarding new engineers
on: user-created
if: has-user-attribute(department=engineering)
steps:
  - uses: notify-user
    with:
      subject: Welcome to Acme Corp Engineering!
      message: Welcome aboard!
  - uses: grant-role
    with:
      role: developer
  - uses: join-group
    with:
      group: /Engineering
'
```

Auto-revoke temporary admin access after 8 days, skipping users who hold the `permanent-admin` attribute, to enforce least-privilege without manual follow-up.

```yaml
name: Temporary admin auto-removal
on: user-role-granted(admin)
if: not has-user-attribute(permanent-admin=true)
steps:
  - uses: notify-user
    after: 7d
    with:
      subject: Admin access expiring soon
      message: Your temporary admin access will be revoked in 24
hours.
  - uses: revoke-role
    after: 1d
    with:
      role: admin
  - uses: notify-user
    with:
      subject: Admin access revoked
      message: Your temporary admin access has been revoked.
```

Speed up the step-runner poll interval during development so short `after` delays (e.g., `after: 30s`) are observed immediately rather than waiting the default 12-hour cycle.

```
$ ./kc.sh start-dev --spi-events-listener--workflow-event-listener--step-runner-task-interval=5s
```

— Names endpoint, flag, all steps, conditions and examples.

Keycloak Workflows in Practice, Part 1:…

THINNER COVERAGE BELOW

02 **Workflows management UI in Admin Console** Not checked yet NEW

45

Workflows can be created and managed through the Admin Console under a new Workflows sidebar entry, which includes a YAML editor for pasting or directly authoring workflow definitions.

— Names the UI path but no deeper mechanism given.

Keycloak Workflows in Practice, Part 1:…

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ Ory Network, Ory Kratos v26.4.23 released NOTES ↗

RANK

Discover the latest features, improvements, and updates to Ory.

> Ory Network's Console gained a redesigned relationships page and a new beta visual diagram of your permission model.

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

01 Visual permission model diagram in Console

Not checked yet NEW

50

The Permissions pages in the Ory Console now include a redesigned relationships page and a new visual diagram of your permission model, currently in beta. You can view it by going to Permissions > Relationships in the Ory Console.

Inspect and visualise your permission model during setup or auditing — no config required.

> 📍 In the Ory Console, go to Permissions › Relationships to see the redesigned relationships page and the new visual permission model diagram BETA.

— UI navigation given but no mechanism or scope detail

Ory Network, Ory Kratos v26.4.23 releas…

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password SaaS Manager #2889 (2026-11-25)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password SaaS Manager adds Claude spend-limit management, exposes roles on single-application API calls, and ships several workflow and UI refinements across SCIM provisioning, support-access auditing, access-level visibility, and budget editing.

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

01 **Roles returned on single-application API calls** Not checked yet

IMPROVED 75

The SaaS Manager API now returns roles — using the same role names as the list endpoint — when getting, creating, or updating a single application, e.g. `GET /api/v1/applications/<app_id>`.

Retrieve a single application's roles via the SaaS Manager API to verify role assignments without needing to call the full applications list.

```
$ curl -X GET
'https://saasmanager.1password.com/api/v1/applications/<app_id>'
\
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/json'
```

— Named endpoint with runnable curl example 1Password SaaS Manager #2889 (2026-11-2…

02 **Immediate audit logging of support-access removal** Not checked yet

IMPROVED 69

Removing support access now writes the removed agents to the audit log and removes them from People immediately, so admins can verify removal in real time via Settings > Support access and the Audit log.

Audit who lost support access after revoking it — the agents now appear in the audit log immediately so you can verify the removal in real time.

> 📍 In SaaS Manager, go to Settings › Support access and turn off support access, then open Audit log to confirm the removed agents are listed.

— Clear UI steps to verify behavior provided 1Password SaaS Manager #2889 (2026-11-2…

03 **Claude spend-limit management in SaaS Manager** Not checked yet NEW

66

New workflow actions let admins set a monthly spend limit for a person's Anthropic Claude usage or clear an existing override, and SaaS Manager now imports each person's Claude spend limit. This requires Claude Enterprise and an API key with spend-limit access.

— Clear requirements and mechanism but no exact UI steps given

1Password SaaS Manager #2889 (2026-11-2…

THINNER COVERAGE BELOW

04 **Access levels and policies counts in applications list**

Not checked yet IMPROVED

55

The Access levels and Access policies columns in the applications list now show linked counts (such as 3 levels and 1 policy), display No policies when an app has none, and apps with no access levels now show an add button to create one.

— UI change described with exact labels but no navigation path

1Password SaaS Manager #2889 (2026-11-2…

05 **SCIM provisioning reuses existing profiles** Not checked yet

IMPROVED

54

SCIM provisioning now reuses an existing SaaS Manager profile sourced from 1Password Enterprise Password Manager when the email domain belongs to your organization, instead of creating a duplicate. Profiles already duplicated before this change are not merged automatically.

— Mechanism and caveat explained but no actionable steps

1Password SaaS Manager #2889 (2026-11-2…

06 **Editing old budgets** Not checked yet IMPROVED

38

Budgets with a start date more than a year ago can now have their value, notes, and owners edited, lifting a prior restriction.

— Brief statement of change with no further detail

1Password SaaS Manager #2889 (2026-11-2…

WAS THIS USEFUL?  

◆ Secure Access (VPN / SASE / ZTNA)

Cato Networks Feedback ⓘ

Zero Trust Isn’t Enough for Agent Security (2026-11-26) BLOG ↗ RANK

Cato Networks is a cloud-native SASE platform that converges networking and security into a single global service for enterprises.

> Cato Networks added AgentSense, a runtime enforcement capability that detects and blocks unsafe AI agent actions before they execute.

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

01 **Untrusted Influence Detection in AgentSense** Not checked yet NEW

42

AgentSense adds Untrusted Influence Detection, which separates trusted instructions from untrusted content and traces how that content influences agent behavior and outputs. This stops indirect prompt injection attacks even when they do not match known patterns, enforced at runtime before unsafe agent actions execute.

– Explains mechanism but no config, API or command details given.

Zero Trust Isn’t Enough for Agent Secur…

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

Tenchi Security Zanshin

Feedback ⓘ

Novidades do Zanshin | Novembro de 2026 NOTES ↗ RANK

Zanshin: a plataforma SaaS que conecta empresas às suas cadeias de suprimentos, promove a transparência e gerencia riscos de segurança cibernética e privacidade.

> Zanshin's November release adds a dedicated Threat Intelligence section for leaked credentials, AI posture checks across AWS, GCP, Azure and OCI, new Azion Web Platform and Jamf Pro scan targets, and a rearchitected ASM v2.0 with distributed parallel scanning and automatic domain discovery.

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

01 Threat Intelligence section for leaked credentials

Not checked yet

NEW 70

Zanshin adds a Threat Intelligence (CTI) section for leaked credentials, accessible under SEGUIDAS > Inteligência de Ameaças for primary organisations or MINHA ORGANIZAÇÃO > Inteligência de Ameaças for third parties, showing detection date, severity, finding type, affected organisation, business impact, and remediation status in a consolidated view.

Review leaked-credential findings for all followed third parties in one place, triaging by severity before the grace period expires

📍 In the left-hand menu, go to SEGUIDAS › Inteligência de Ameaças. Use the severity and risk-level summary cards (critical / high / medium) to prioritise credential findings across followed organisations.

— Exact UI paths and displayed fields named, no API

Novidades do Zanshin | Novembro de 2026

02 New scan targets: Azion Web Platform and Jamf Pro

Not checked yet

NEW 65

Zanshin adds scan target support for Azion Web Platform (found under IaaS | PaaS on the Environments Panel) and Jamf Pro (found under Ferramentas de Segurança on the Environments Panel).

Add Jamf Pro as a continuously monitored scan target to bring mobile-device management into your attack-surface scope

📍 On the Painel de Ambientes page, locate the Ferramentas de Segurança

section and select Jamf Pro to register it as a new scan environment.

— Exact UI location named for both new targets Novidades do Zanshin | Novembro de 2026

THINNER COVERAGE BELOW

03 **ASM v2.0 distributed parallel scanning** Not checked yet IMPROVED

55

ASM v2.0 replaces sequential execution limits with a distributed parallel-scanning architecture, enabling full coverage of large attack surfaces without timeout interruptions or dropped endpoints.

— Explains mechanism and before/after, no config surface

Novidades do Zanshin | Novembro de 2026

04 **Automatic discovery of unmapped domains** Not checked yet NEW

55

Zanshin now automatically discovers regional, legacy, or brand-related domains associated with an organisation that were not previously mapped, even if never manually registered, and sends automated notifications whenever these newly discovered domains are identified and added to the active scan scope.

— Explains mechanism and notification behaviour, no UI path given

Novidades do Zanshin | Novembro de 2026

05 **SAQ reference attachments and requester branding** Not checked yet

IMPROVED

55

Assessment questionnaires (SAQ) now support attaching reference documents — such as instructions or guidelines — when creating or editing a questionnaire to help respondents understand and answer questions, and the Assessment header now shows the requesting organisation's logo so third parties can immediately identify which following organisation sent the assessment and will receive the responses.

Attach reference guidelines to an SAQ so respondents in a third-party assessment know exactly what evidence to supply

📍 In the SAQ editor, open an existing questionnaire or create a new one, then use the file-attachment option to upload instruction documents (e.g. a PDF policy guide) before sending the assessment.

— UI steps given for attaching documents, no config key

Novidades do Zanshin | Novembro de 2026

06 **Dynamic score updates with remediation grace period**

Not checked yet IMPROVED

50

Score calculation now updates dynamically as remediation tasks are completed and verified, giving immediate credit rather than waiting for prolonged calculation cycles, and newly discovered findings now include a grace period before they affect the overall score, giving security and engineering teams time to triage, prioritise, and remediate emerging risks before scores fluctuate.

— Explains behaviour change but no exact timing or thresholds

Novidades do Zanshin | Novembro de 2026

---

07 **AI posture checks for major clouds** Not checked yet NEW

30

Zanshin now includes AI posture checks for AWS, GCP, Azure, and OCI.

— Names providers but no mechanism or check detail

Novidades do Zanshin | Novembro de 2026

---

08 **Assessment reports in Brazilian Portuguese** Not checked yet NEW

25

Assessment reports are now available in Brazilian Portuguese (`pt-br`).

— One-line locale addition with no further detail Novidades do Zanshin | Novembro de 2026

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

LogicGate

Feedback ⓘ

ISF World Congress Milan 2026: LogicGate Recap (2026-11-27) BLOG ↗ RANK

LogicGate Risk Cloud is the leading AI-powered GRC platform with 30+ applications for enterprise risk, third-party risk, cyber risk and compliance management.

> LogicGate rolled out GRC Agents to general availability across four risk domains and introduced the Insights Agent for plain-language risk queries.

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

01 **GRC Agents reach general availability with Insights Agent**

Not checked yet NEW 44

GRC Agents are now generally available for third-party risk management, enterprise risk management, AI governance, and business continuity management. The newly released Insights Agent, one of these GRC Agents, lets users ask about risk in plain language as part of LogicGate's move toward orchestrated, autonomous GRC.

— Names four domains but no commands, UI path, or mechanism detail.

ISF World Congress Milan 2026: LogicGat…

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 25th, 2026 NOTES ↗ RANK

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

> Compyl added an organization-wide setting to hide the summary charts above the grid in Risk and Third Party modules.

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

01 **Hide charts in Risk and Third Party grids** Not checked yet NEW

55

Admins can turn off Show charts above the grid separately in Risk Settings and Third Party Settings to hide the charts above the grid for everyone in the organization.

. . . .

— Names exact setting and location but no further mechanism. Nov 25th, 2026

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.9.0 (2026-11-28) BLOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer 6.9.0 adds agent verification badges and Slack rich-card output, masks deployment credentials with a new API validation mode, and hardens AI Watch hook device identity alongside audit, policy and Kubernetes operator fixes.

└──▷ WHAT SHIPPED · 11 FEATURES most completely described first ? what's the number?

01 Deployment YAML credential masking and placeholder validation

Not checked yet

95

BREAKING

Deployment YAML credentials in env values longer than the exact strings `$$DEPLOYMENT_URL`, `$$RUNLAYER_URL` or `$$RUNLAYER_OAUTH_CALLBACK_URL` are now masked in `runlayer deploy pull`, the YAML editor and the export API; `PUT /api/v1/deployments/{id}` with `yaml_content` preserves the stored value when a `***` placeholder is submitted and returns a `422` when that key has nothing stored.

Pull a deployment's YAML after upgrading to 6.9.0 to see masked credential placeholders, then restore the real values before redeploying.

```
$ runlayer deploy pull > deployment.yaml
# Edit deployment.yaml: replace each ${VAR} placeholder with its
real credential value
runlayer deploy deployment.yaml
```

Save a deployment via the API when a masked credential is already stored — send `***` to preserve the existing value rather than overwriting it.

```
$ curl -X PUT https://your-runlayer-host/api/v1/deployments/{id}
\
  -H 'Content-Type: application/json' \
  -d '{"yaml_content": "env:\n  RUNLAYER_URL: \"***\""}'
```

— Exact CLI command, endpoint, method and placeholder behavior given. 6.9.0 (2026-11-28)

02 AI Watch hook identity and health checks hardened

Not checked yet

90

BREAKING

Hooks on devices using the managed organization key now identify calls solely by the device's hardware ID, so a device named in the hook payload or in `RUNLAYER_DEVICE_ID` can no longer borrow another device's user exemptions or have calls attributed to that device's user; if the hardware ID cannot be read, Protect and Enforce block MCP and local tool calls rather than sending them without a device. Hook health now covers every VS Code profile found on a device — VS Code Insiders, Exploration and VSCodium — not only the default stable profile, reporting the device as drifted if any of those user settings would not load the hooks or have `chat.useHooks` turned off. When AI Watch cannot identify an MCP server because a tool call's name matches more than one configured server, the audit log entry now explains that more than one server fits the name.

— Names exact env var and settings but no runnable command. 6.9.0 (2026-11-28)

---

03 **Kubernetes operator sandbox NetworkPolicy ownership**

Not checked yet

70

BREAKING

The Kubernetes operator now names each `RunlayerInstance`'s sandbox `NetworkPolicy` resources uniquely, refuses to modify policies written by another instance, and rejects a second same-named instance in a shared agents namespace with `SandboxNetworkPolicyNotOwned`.

— Names exact error code but gives no remediation steps. 6.9.0 (2026-11-28)

---

04 **Microsoft 365 Copilot shared client scoped to owning org**

Not checked yet

70

BREAKING

Microsoft 365 Copilot connections made through a shared Copilot client with no organization recorded now end at the next refresh for users outside the organization Runlayer treats as the client's owner; affected deployments with more than one organization must set up a dedicated Copilot client and enter its client ID and secret in Copilot to reconnect.

— Gives remediation steps but no exact UI path or API. 6.9.0 (2026-11-28)

---

05 **Verified agents with manager sign-off** Not checked yet NEW

65

Agent managers can now mark agents as verified, the same way they verify skills; verified agents show a check on agent cards and the agents list gains a filter for them. Editing the agent's name, instructions, model, connectors or skills removes the mark until a manager re-verifies.

— Explains verification mechanism and revocation trigger, only UI path given.

6.9.0 (2026-11-28)

---

06 **Audit event triggers scoped to creator's visible events**

Not checked yet

BREAKING

65

An Audit event trigger now starts a run only for events its creator can see in the Audit log; without View org audit logs permission that is limited to the creator's own events, and existing triggers narrow automatically to this scope.

— Explains permission scoping but names no config surface. 6.9.0 (2026-11-28)

THINNER COVERAGE BELOW

07 **Sort agent sessions by cost** Not checked yet NEW

55

An agent's Sessions list can now sort by cost with the Most expensive option under Sort by in the Filter menu, surfacing the priciest runs first.

Review the most expensive agent runs first to identify costly sessions for cost governance.

📍 In the Agents section, open the Sessions list for an agent, open the Filter menu, and choose Most expensive under Sort by .

— Exact UI navigation given but limited mechanism detail. 6.9.0 (2026-11-28)

08 **OAuth tokens invalidated on directory-sync org move**

Not checked yet

BREAKING

55

OAuth access tokens issued to a client belonging to one organization stop working for users who have been moved to a different organization via directory sync; refresh attempts are refused and the refresh token is revoked, requiring the user to reconnect through a client their current organization owns.

— Explains behavior but no command or config to remediate. 6.9.0 (2026-11-28)

09 **Slack rich-card output for agent messages** Not checked yet NEW

50

Agents that post to Slack can now attach a native Slack block card under their message containing tables, bar, line and pie charts, key-value fields and link buttons.

— Lists supported card elements but no config to enable it. 6.9.0 (2026-11-28)

10 **Deny Policy Enhancements** Not checked yet IMPROVED

50

A deny policy on an agent account for session content now also withholds run content from that agent when it asks the Runlayer MCP to suggest eval runs while acting for a person, and revoking a user's direct connector access now only removes allow policies and no longer deletes or narrows deny policies, with the confirmation indicating if a group or Everyone policy still grants access.

— States behavior but gives no config key or command. Describes behavior change without a config or API surface.

6.9.0 (2026-11-28)

11 Completed agent approvals view in Inbox

Not checked yet NEW

50

Completed agent approvals in Inbox now show approved, denied and expired requests from the last 7 days when Show completed is turned on, including who decided and when.

— Exact UI toggle named but shallow mechanism. 6.9.0 (2026-11-28)

└──▷ BREAKING ON UPGRADE

! Deployment YAML env values longer than the exact strings `$$DEPLOYMENT_URL`, `$$RUNLAYER_URL` or `$$RUNLAYER_OAUTH_CALLBACK_URL` — such as `$$RUNLAYER_URL?api_key=...` — are now masked in `runlayer deploy pull` output and the export API; a file pulled with `runlayer deploy pull` now contains a `${VAR}` placeholder, requiring you to write the real value back before running `runlayer deploy`. Rotate any credential that was stored in such a value.

! `PUT /api/v1/deployments/{id}` with `yaml_content` now returns a `422` when a masked `***` placeholder is submitted for a key that has no stored value, rather than saving a literal `***`.

! Microsoft 365 Copilot connections made through a shared Copilot client with no organization recorded now end at the next refresh for users outside the organization Runlayer treats as the client's owner; affected deployments with more than one organization must set up a dedicated Copilot client and enter its client ID and secret in Copilot to reconnect.

! OAuth access tokens issued to a client belonging to one organization stop working for users who have been moved to a different organization via directory sync; refresh attempts are refused and the refresh token is revoked, requiring the user to reconnect through a client their current organization owns.

! Existing Audit event triggers that previously fired on events the creator could not see in the Audit log now narrow automatically and go quiet for those events.

! The Kubernetes operator renames sandbox `NetworkPolicy` resources to include a per-instance qualifier; a second `RunlayerInstance` with the same name in a different namespace sharing one agents namespace now fails with `SandboxNetworkPolicyNotOwned` rather than overwriting the first instance's policies.

! AI Watch hooks on devices using the managed organization key no longer accept the device identity from the hook payload or from `RUNLAYER_DEVICE_ID`; if the hardware ID cannot be read, Protect and Enforce block MCP and local tool calls rather than passing them through.

WAS THIS USEFUL? 👍 👎

EVALUATE

◆ AI Observability & Evals

ongrid

Feedback ⓘ v0.22.1 NOTES ↗ RANK

> ongrid now lets teams configure automatic knowledge sync on a per-repository basis.

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

01 **Per-repository automatic knowledge sync configuration**

Not checked yet NEW 31

Automatic knowledge sync can now be configured on a per-repository basis, rather than applying uniformly across all repositories.

— No config key, UI path, or mechanism given beyond prose description v0.22.1

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Zenity

Feedback ⓘ changelog · Nov 28, 2026 CHANGELOG ↗ RANK

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

> Zenity ships MCP Gateway as a generally available control point for governing every MCP tool call.

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

01 **MCP Gateway reaches general availability** Not checked yet NEW

33

Zenity's MCP Gateway is now generally available, serving as a single control point for governing every MCP tool call.

— Names the capability but no mechanism or config details changelog-20261128-3c1255b3

WAS THIS USEFUL? 👍 👎

◆ Model & Agent Protection

h5i

Feedback ⓘ v0.5.7 NOTES ↗ RANK

> h5i's v0.5.7 release introduces an installable CLI plugin system, adding browser automation with a policy-checking proxy, a websec workbench for request replay, a recon crawler for attack surface discovery, and security regression testing with OpenAPI coverage gating.

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

01 **Security regression testing with OpenAPI coverage gating**

Not checked yet NEW 92

`h5i test .h5i-tests/tests --target http://localhost:3000` runs repository-owned attack flows and reports results in JSON and JUnit output; supplying `--min-coverage` (e.g. `h5i test --target https://staging.example.com --openapi openapi.yaml --min-coverage 80`) enforces a minimum OpenAPI coverage threshold.

— Full command, flags, and output formats named. v0.5.7

02 **Browser automation with policy-checking proxy** Not checked yet NEW 88

`h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` open a Chromium session, capture every request through a policy-checking proxy, and let you review the full session afterwards.

Watch every HTTP request a Chromium session makes — and what was refused — before any bytes leave the proxy.

```
$ h5i browser open https://example.com
h5i browser requests
h5i browser audit
```

— Names all four subcommands and the proxy mechanism. v0.5.7

03 **Websec workbench plugin for request replay** Not checked yet NEW 86

Install the plugin with `h5i plugin install websec --from ./h5i-websec` to replay and modify captured HTTP requests, e.g. `h5i websec replay req_42 --set query.id=456`, and inspect session traffic with `h5i websec requests`.

— Exact install command and replay syntax provided. v0.5.7

04 **Recon plugin for attack surface discovery** Not checked yet NEW

86

Install the plugin with `h5i plugin install recon --from ./h5i-recon` to record what a target exposes and where each discovered candidate came from, using `h5i recon crawl --max-requests 200`.

Discover attack surface on a target by crawling up to 200 requests and recording every found candidate with its source.

```
$ h5i plugin install recon --from ./h5i-recon
h5i recon crawl --max-requests 200
```

— Install and crawl commands with flag given. v0.5.7

WAS THIS USEFUL? 👍 👎

ToolHive

Feedback ⓘ v0.70.0 NOTES ↗ RANK

> ToolHive v0.70.0 strengthens the vMCP access-control model with call-bound tokens and per-user upstream credential storage, while adding readiness checks, OpenShift compatibility, and call-chain resilience improvements.

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

01 **Call-bound access tokens at the vMCP** Not checked yet NEW

65

ToolHive now enforces call-bound access tokens at the vMCP for both tool calls and tool-listing operations, so each call carries its own scoped token rather than a shared session credential. `tool:call` now also implies `tool:read` for tool listing, so a principal granted call access no longer needs an explicit read grant.

– Names the permission model and exact scopes but no config example v0.70.0

02 **OpenShift /finalizers subresource support** Not checked yet NEW

60

Adds a `/finalizers` subresource for OKD/OpenShift installs to satisfy OpenShift's RBAC requirements for owner-reference cleanup.

– Names the exact subresource and platform requirement v0.70.0

THINNER COVERAGE BELOW

03 **Cedar policy access to ID token claims** Not checked yet NEW

45

Cedar can now read user claims the upstream IdP asserts only in its ID token, expanding the attributes available for policy evaluation.

– Names Cedar and ID token source but no policy syntax given v0.70.0

04 **MCP initialize readiness check** Not checked yet NEW 45

Adds an MCP initialize readiness check so ToolHive waits for a backend to complete the MCP handshake before marking it ready.

– Clear behaviour described, no flag or config name given v0.70.0

05 **Per-user upstream credential storage** Not checked yet NEW

40

Adds per-user upstream credential reads and user-keyed upstream token storage, so each user's upstream credentials are stored and retrieved independently.

– Describes the mechanism but no API or config detail v0.70.0

06 **Cancellation forwarding to vMCP backends** Not checked yet IMPROVED

40

Forwards a client's cancellation to vMCP backends, so cancellation signals propagate through the full call chain.

— Explains propagation mechanism, no interface detail given v0.70.0

07 **Customizable MCP instructions field** Not checked yet NEW

35

Operators can now set the instructions MCP endpoints return, allowing per-deployment customization of the instructions field.

— States the capability but not the config key or field path v0.70.0

08 **Non-fatal secondary chain legs** Not checked yet IMPROVED

35

Secondary chain legs can now be made optionally non-fatal, so a failure in a non-critical chain step does not abort the whole request.

— Describes behaviour change with no config or flag named v0.70.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Bugcrowd Researcher Platform Feedback ⓘ

Request a Response (RaR) Reasons Update (2026-11-26) BLOG ↗ RANK

Bugcrowd Researcher Platform is a crowdsourced security marketplace where ethical hackers find and report vulnerabilities in client programs to earn rewards.

> Bugcrowd Researcher Platform refined its Request a Response flow with recipient-aware reason categories and submission guidance.

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

01 **Audience-specific Request a Response reasons** Not checked yet

IMPROVED 45

Request a Response now offers reason categories matched to the recipient (Bugcrowd or Customer, including self-triaged programs), with short guidance under each reason and evidence hints shown in the details field.

— Describes behaviour but no exact UI path or limits

Request a Response (RaR) Reasons Update…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 28, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue