THE CYBER TOOLCHAIN

NO. 149

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 24, 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 145 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 57 tools

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

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. Velociraptor can now seed endpoint tool inventories over SMB with hash pinning, so air-gapped DFIR

vim ? help

collections stop failing on missing downloads. Prowler reaches Microsoft 365 tenants in sovereign and air-gapped clouds. Auth0's Poppy beta gives a personal AI agent scoped, revocable access instead of the user's password.

RESPOND

Run Velociraptor collections in an air-gapped estate without hand-staging binaries

Third-party tools can be distributed over SMB and other URL schemes with hash pinning, and a built-in artifact generates a downloader script to pre-populate the tool inventory. Responders stop manually copying sysinternals-style helpers to isolated hosts and stop losing a collection midway because an endpoint couldn't reach the internet.

Velociraptor

PROTECT

Hand an AI agent revocable access instead of your password

Poppy beta lets a personal agent hold scoped, revocable delegated access. The current practice — pasting a user's own credentials or a long-lived token into an agent — leaves no way to see what the agent did or cut it off without a password reset.

Auth0

PROTECT

Stop a malicious package landing before the cooldown window expires

Secure Registry now exempts security fixes drawn from multiple vulnerability databases from the quarantine cooldown, and adds beta typosquatting protection for PyPI and Maven. Cooldowns normally delay legitimate patches as the price of blocking fresh malicious releases; the exemption removes that trade-off.

StepSecurity

IDENTIFY

Assess Microsoft 365 posture in tenants a scanner couldn't previously reach

Prowler 5.52.0 scans M365 in sovereign and air-gapped clouds and widens IPv6 exposure detection; Plerion adds four Azure Storage checks and reworks its S3 public-access logic. Sovereign tenants have been a

> reporting blind spot that teams covered with spreadsheets and manual console review.
>
> Prowler · Plerion

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Password & Credential Cracking

brutespray

Feedback ⓘ v2.7.7 NOTES ↗ RANK

Brutespray is a fast, multi-protocol credential brute-forcer. Parses Nmap, Nessus, and Nexpose output to automatically test default and custom credentials across 30+ protocols.

> BruteSpray v2.7.7 adds detection of any-password-accepting devices with a confirmation step and tightens how connection failures are counted before the circuit breaker trips.

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

01 Any-password hit confirmation for brute-force results

Not checked yet

82

BREAKING

BruteSpray now detects devices that accept any password (DVRs, IP cameras, managed switches) by re-checking each apparent hit with a random password; demoted hits are written as `WARN any-password` findings instead of success-file entries. The re-probe can be skipped with `-no-confirm-hits` for environments where the double-check would cause account lockouts or excessive noise. Existing success files from RDP/HTTP runs should be re-checked, since previously logged successes may now be reclassified under this pattern.

Skip the random-password re-probe entirely when targeting environments where the double-check would cause account lockouts or excessive noise.

```
$ brutespray -f nmap.xml -u admin -p passlist.txt -no-confirm-hits
```

▶ View command output v2.7.7

— Names mechanism, flag and output change, but no numeric details v2.7.7

02 Per-host connection failure cap enforcement

Not checked yet

68

BREAKING

`-max-conn-fails n` now caps connection errors per host before the circuit breaker trips, enforced atomically before dialing regardless of thread count, so a dead host never logs more errors than the configured limit. This changes prior behavior: the flag now counts only connection errors (not total attempts), so hosts that previously tripped the breaker on auth failures may no longer trip it.

— Names flag and behavior change, no usage example given v2.7.7

└──▷ BREAKING ON UPGRADE

! `-max-conn-fails n` now counts only connection errors per host (not total attempts); hosts that previously tripped the breaker on auth failures may no longer trip it.

! Hosts that previously logged successful credentials but accept any password now appear as `WARN any-password` findings rather than success-file entries; existing success files from RDP/HTTP runs should be re-checked for this pattern.

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

AIPentest CyberStrikeAI

Feedback ⓘ v1.8.1 NOTES ↗

RANK

CyberStrikeAI is an ai-native penetration testing platform in Go that orchestrates offensive tools over MCP and tracks runs in a web dashboard.

> CyberStrikeAI migrated its agent backend to a new Pi engine with several reliability and model-handling refinements, and expanded its vulnerability workflow with grouping, retesting, and improved list display.

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

01 **Migration to Pi agent engine** Not checked yet NEW 60

CyberStrikeAI now uses the Pi agent engine, replacing the previous engine. The Pi engine can omit the `store` field on a per-channel basis for compatible gateways, preserves the system role when working with compatible reasoning models, reports incomplete runs and model stream failures with clear error messages, improves session management for more reliable multi-turn interactions, and optimizes DeepSeek input handling and local model context adaptation.

– Names engine, DeepSeek and store field but no flags v1.8.0 v1.8.1

THINNER COVERAGE BELOW

02 **Vulnerability list timestamps, sorting, and pagination**

Not checked yet IMPROVED 45

The vulnerability list now shows timestamps, supports oldest-first sorting, and has independent pagination separate from other list views.

– Concrete UI changes but no config or endpoint named v1.8.1

03 **Grouped vulnerability browsing by category** Not checked yet NEW 36

Vulnerabilities can now be viewed in a browsable grouped view so practitioners can triage findings by category.

– Describes behavior but no UI path or filter names v1.8.0

04 **Retest tracking for vulnerabilities** Not checked yet NEW

36

CyberStrikeAI adds retest tracking to follow the remediation status of individual vulnerabilities.

— Clear purpose but no mechanism or field names v1.8.0

05 **Token usage shown in history timelines** Not checked yet NEW

28

The frontend now shows token usage in history timelines.

— Brief one-line UI addition with no further detail v1.8.1

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

OpenQodex

Feedback ⓘ openqodex@0.23.0 NOTES ↗ RANK

OpenQodex is a join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> OpenQodex's latest releases add configurable Codex reviewer model and effort settings, plus a major transparency push for the callers/code graph system — a new `unknown.gaps` field, better Python caller detection, named floor causes for framework route handlers, a golangci-lint correctness fix, and faster graph builds.

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

01 **Configurable Codex reviewer model and effort** Not checked yet

NEW 95

Set `reviewer_model` (a model ID) and `reviewer_effort` (`low`, `medium`, `high`, or `xhigh`) in `~/.openqodex/config.yaml` to control which Codex model and reasoning level is used for every review. Pass `--reviewer-model` and `--reviewer-effort` to the `review` command to override these for a single run without touching persistent config, and check which model/effort are active via `openqodex doctor`, which now names both in its output table.

Pin the Codex reviewer to a specific model and effort level for all future reviews on this machine.

```yaml
reviewer_model: gpt-6.1-sol
reviewer_effort: high
```

Run a one-off review at maximum reasoning effort without changing your persistent config — useful for a final pre-merge review on a large diff.

```
$ openqodex review --reviewer-model gpt-6.1-sol --reviewer-effort xhigh
```

— Names exact config keys, flags, values and doctor check openqodex@0.23.0

02 **`unknown.gaps` field in callers and callees answers**

Not checked yet NEW 85

A new `unknown.gaps` field in `callers` and `callees` answers lists every place the graph gave up — with cause, file, line, and what to check next — covering syntax errors and unreadable regions across the seven supported languages, and runtime reflective lookups (`getMethod` , `getDeclaredMethod` , `MethodByName` , `getattr`). Multi-hop `callers` and second-hop `impact` answers now say may be incomplete and name the specific caller with a gap when any walked caller has a graph or framework gap.

— Names the field, its contents and propagation, but no runnable example `openqodex@0.22.2`

---

03 **Named floor causes for framework route handlers in `callers`**

Not checked yet IMPROVED 85

`callers` now floors with a named cause instead of silently dropping handlers or showing an empty list as complete: Express routers mounted under a computed prefix (such as app.use(process.env.PREFIX, router)) report routes as unknown with a computed path; FastAPI and Nest handlers dropped by plugin limits (long dotted names, deep parentheses, guards past the 64th, classes with over 1,000 members); framework handlers with computed route paths (FastAPI f-string paths, Go constants over the string limit, Express computed paths) each get a named cause; and Next.js and Nest route-handler files over 256 KB that the plugins skipped now name the unread file.

— Enumerates every named limit and framework case with specific thresholds

`openqodex@0.22.2`

---

04 **Expanded Python caller detection in `graph callers`**

Not checked yet IMPROVED 80

`graph callers` now recognizes two more Python usage patterns so callers floor correctly: a function used as a bare decorator (`@wrap` or `@pkg.wrap` , including on methods), not only @wrap() calls, and a function passed inside extra parentheses (such as Depends(((dep)))), which now counts as used as a value.

— Exact syntax patterns and before/after given, no flag to run `openqodex@0.22.2`

---

05 **golangci-lint failure detection fix** Not checked yet IMPROVED 60

golangci-lint now reads as failed when any Go module in the change cannot type-check, even when another module produced findings, so unchecked packages are no longer silently hidden.

— Clear before/after but no command or config shown `openqodex@0.22.2`

THINNER COVERAGE BELOW

06 **Faster code graph builds** Not checked yet IMPROVED 40

The code graph now builds faster on deeply nested code, barrels with circular `export *` paths, and files with many `eval` or global-object sites.

— Named scenarios but no numbers or way to invoke/verify openqodex@0.22.2

WAS THIS USEFUL? 👍 👎

Vercel deepsec

Feedback ⓘ changes since 2026-11-19 CODE ↗

Deepsec is a security harness for finding vulnerabilities in your codebase powered by coding agents

deepsec adds a `--spawn-rate` limit for sandbox runs and now retries content-filter-blocked batches in smaller groups instead of failing them

└──▷ **TRY IT**

Throttle sandbox startup to avoid connection storms when launching many workers at once against a large codebase.

```
$ deepsec sandbox process --spawn-rate 5
```

Control how fast sandboxes start during a `sandbox process` run with the new `--spawn-rate` flag, which defaults to 10 per minute. Not checked yet

deepsec now retries content-filter-blocked batches in smaller groups rather than failing every file in the batch, reducing unanalyzed files from ~51–61% of batches to roughly 0.28% in a 31,600-file run. Not checked yet

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Issue statuses and repository boards are in public preview (2026-11-23) BLOG ↗

RANK

GitHub is a web-based platform for version control and collaborative software development using Git repositories.

> GitHub rolled out issue statuses and board views to repository Issues (with API, automation-trigger, and Projects-sync support) in public preview, while tightening code-coverage upload validation and adding a user_id column to usage reports for tracking deleted accounts.

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

01 **Issue statuses with board views and API support** Not checked yet

NEW 88

Repository issues now carry a status field filterable with `status:` across the repository issues page, the issues dashboard, and repository boards. Status can be read, set, or assigned via GraphQL, REST API, and the MCP server, letting automations and agents manage it programmatically, and status changes can trigger Copilot automations. Saved views can be switched to a board layout grouped by status with drag-to-update columns, swimlanes, visible fields, and sorting, and any issue field can now be added via the cog menu in the issue sidebar regardless of issue type.

▶ Watch video

Find all issues currently blocked in review on a repository issues page or board.

```
$ status:In review
```

Switch a repository saved view to the board layout, grouped by status, to drag issues across columns and update their status.

📍 1. Open any saved view in your repository's Issues tab. 2. Click the View options menu (⚙️). 3. Select Board as the layout. 4. Set Group by to Status. 5. Drag an issue card to a different column to update its status.

— Names filter syntax, APIs, and UI steps for board layout

Issue statuses and repository boards ar…

02 **Cobertura format validation for coverage uploads** Not checked yet

80

BREAKING

The `upload-code-coverage` action now fails fast with the error 'This file is not a Cobertura XML report. Upload a Cobertura XML report.' when the uploaded report isn't Cobertura XML, replacing silent no-op uploads that previously showed no coverage on pull requests. Valid Cobertura reports are now processed more leniently: individual lines missing a usable line number or hit count are skipped rather than discarding the whole report.

— Names the action and exact error message, explains before/after

Code coverage uploads now flag invalid …

---

03 **Automated issue status transitions from PR and assignee activity**

Not checked yet NEW

73

Linking a draft PR moves an issue to In progress, marking that draft PR ready moves it to In review, linking an already-open PR moves it to In review, assigning a user moves it to Todo, and assigning an agent moves it to In progress.

— Exact transition rules given but no command to invoke

Issue statuses and repository boards ar…

---

04 **user_id column added to usage reports** Not checked yet BREAKING

69

A new `user_id` column now appears in the detailed usage report, the AI usage report, and the premium request usage report — including REST API exports — enabling tracking of usage from deleted accounts and users across username changes. Every column after `username` shifts one position right in these reports starting December 1, 2026, so tools reading by column position rather than header name will break.

— Names exact column, reports, and migration deadline

Usage reports will include a user ID co…

THINNER COVERAGE BELOW

---

05 **Bidirectional sync between issue status and Projects**

Not checked yet NEW

45

When a Project has a status value matching an issue's status, changing either the Project's status or the issue's status updates the other.

— States mechanism but no config or steps shown Issue statuses and repository boards ar…

└──▷ BREAKING ON UPGRADE

! Workflows using `upload-code-coverage` with a non-Cobertura XML report will now fail with the error 'This file is not a Cobertura XML report. Upload a Cobertura XML report.' — uploads

that previously appeared to succeed will now break.

! Every column after `username` shifts one position to the right in the detailed usage report, the AI usage report, and the premium request usage report (including REST API exports) — scripts or tools that read these reports by column position will break unless updated to read columns by header name before December 1, 2026.

WAS THIS USEFUL? 👍 👎

Snyk Code

Feedback ⓘ changelog · Nov 23, 2026 BLOG ↗ RANK

Snyk Code is a static application security testing (SAST) tool that scans developer code in real time to identify and fix security vulnerabilities.

> Snyk Code fixed SCM import scanning so that `.snyk` exclusion files saved in Latin-1 or Windows-1252 encodings are now correctly applied.

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

01 **Encoding support for .snyk exclusion files** Not checked yet

IMPROVED 43

A `.snyk` file saved in Latin-1 or Windows-1252 encoding now correctly applies its exclusion rules during SCM imports.

— Names exact encodings and file but no scope of impact 2026-11-23-4eafc00e

WAS THIS USEFUL? 👍 👎

skylos

Feedback ⓘ v5.8.0 NOTES ↗ RANK

Skylos is an open source local-first PR scanner that finds dead code, security bugs, secrets, quality regressions, and AI-code mistakes before merge. For first timers refer to https://duriantaco.github.io/skylos/repo-map/

> Skylos added a cloud dashboard with PR-blocking CI gating through a GitHub App, plus UML change reports on PRs, a bundled Semgrep rule pack, new Python and Ruby security checks, stacked-PR verification, and more accurate dead-code detection.

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

01 Cloud dashboard and PR-blocking CI gate

Not checked yet NEW

91

Run `skylos cicd init` to write `.github/workflows/skylos.yml` and get the git commands and GitHub App link needed to enable pull-request blocking. Skylos Cloud uploads scan findings to a hosted dashboard at `skylos.dev/dashboard` without sending source code, and the GitHub App integration authenticates CI scans via OIDC, removing the need to store any token or secret in the repository. Only CI-uploaded scans can block merges or update the project verdict; local scans are saved for review only.

Bootstrap CI gating so that pull requests failing the Skylos check cannot be merged.

```
$ skylos cicd init
```

— Names exact command, workflow file, auth mechanism and verdict rule product docs

02 Bundled Semgrep rule pack via vulnhunter

Not checked yet NEW

73

skylos adds `skylos rules vulnhunter` with a bundled Semgrep rule pack for running Semgrep-backed rules alongside native checks.

Run the bundled Semgrep-backed rule pack against your repo alongside all other checks to catch vulnerabilities Semgrep rules cover.

```
$ skylos rules vulnhunter
```

— Runnable command named but mechanism of rule pack is thin v5.8.0

03 PR-stack verification with skylos done

Not checked yet NEW

70

skylos now checks every layer of a PR stack with `skylos done`, verifying that every commit in a stacked PR series passes the gate before merging, not just the tip branch.

Verify that every commit in a stacked PR series passes the done gate before merging, not just the tip branch.

```
$ skylos done --base main
```

— Runnable command with flag shown, scope of stacking explained v5.8.0

THINNER COVERAGE BELOW

04 **New security checks for Python and Ruby** Not checked yet NEW

55

skylos now detects XML external entity (XXE) vulnerabilities and plain HMAC comparisons in Python code, and checks Rails controllers for SQL injection and SSRF in Ruby code.

— Names vulnerability classes and languages but no command shown v5.8.0

05 **Improved dead-code detection accuracy** Not checked yet IMPROVED

35

skylos now produces more accurate dead-code results for TypeScript, JavaScript, and Rust with less noise.

— Names languages improved but no metric or mechanism given v5.8.0

06 **UML change report on PR scans** Not checked yet NEW

30

skylos now adds a UML change report to pull request scans, giving reviewers a diagram of what the PR changed structurally.

— Describes purpose but no command or output format given v5.8.0

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-20 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds an LLM advisor to the bisect agent to supervise kernel bisection and keep it on track within its iteration budget

An `LLMAdvisor` now supervises the bisect agent, using `aflow.DeepReasoningModel` with `MinIterations: 20` and `EvaluationInterval: 20` to redirect looping behaviour and prompt early termination when the root cause is confirmed.

Not checked yet

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Nov 22, 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 added a new Org Settings API capability to programmatically control which ignore sources an organization accepts.

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

01 **Org Settings API for ignore sources** Not checked yet NEW

69

A new Org Settings API lets admins control which ignore sources are accepted org-wide — app and API only, or also CLI and `.snyk` files — mirroring the existing UI setting under Settings → General.

— Names API and config surface but lacks endpoint path or examples. 2026-11-22-cfcbe2fa

WAS THIS USEFUL? 👍 👎

zizmor

Feedback ⓘ v1.33.0 NOTES ↗ RANK

Zizmor is a static analysis for GitHub Actions (and more)

> zizmor v1.33.0 adds a new audit for dangerous Git checkout patterns and extends several existing audits with semver-granular dependabot cooldowns, workflow lockfile awareness, and updated runner/action recommendations.

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

01 Dependabot.yml audit enhancements

Not checked yet IMPROVED 65

The `dependabot-cooldown` audit now supports `semver-{major,minor,patch}-days` cooldown settings in `dependabot.yml`, and zizmor now recognizes `apm` as a valid `package-ecosystem` value in `dependabot.yml`.

– Names exact config keys but no usage example v1.33.0

THINNER COVERAGE BELOW

02 New dangerous-checkout audit

Not checked yet NEW 55

A new `dangerous-checkout` audit detects dangerous Git checkout patterns in GitHub Actions workflows, starting with coverage for `actions/checkout`.

– Names audit and initial action coverage but no mechanism detail v1.33.0

03 superfluous-actions recommends gh release download

Not checked yet IMPROVED 50

The `superfluous-actions` audit now recommends `gh release download` as a replacement for `robinraju/release-downloader`.

– Names audit, action, and replacement command v1.33.0

04 Workflow lockfile awareness in impostor-commit and ref-confusion

Not checked yet IMPROVED 45

The `impostor-commit` audit is now aware of workflow lockfiles, reducing noise for pinned actions, and the `ref-confusion` audit is now aware of workflow lockfiles as well.

– Names two audits but mechanism of lockfile awareness unclear v1.33.0

05 self-hosted-runner recognizes xcode-27

Not checked yet IMPROVED 30

The `self-hosted-runner` audit now recognises `xcode-27` GitHub runners.

— Bare statement of new runner recognition v1.33.0

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

rearm

Feedback ⓘ 26.11.59 NOTES ↗ RANK

ReARM - Release Governance Platform

> ReARM's latest release makes the programmatic API's nginx proxy timeout configurable and adds design-checkpoint and split-task workflows to agent task boards, alongside smaller notification and statistics tweaks.

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

01 **Design checkpoints and split tasks in agent boards**

Not checked yet NEW 62

Agent task boards now carry a design checkpoint before the build role, with an option to send a task on without a split. They also support split tasks that carry their parent's reviewed design, shown via a header chip, hop row, documents group, and a marker on proposed children.

– Describes new workflow and UI elements but no exact commands or keys. 26.11.59

THINNER COVERAGE BELOW

02 **Configurable nginx proxy timeout for API** Not checked yet NEW 52

The nginx proxy timeout for the programmatic API is now configurable, with a default of 600 seconds, so long-running calls are no longer cut off.

– Names mechanism and default but no exact config key given. 26.11.59

03 **Opt-in board documents in notifications** Not checked yet NEW 28

Board documents can now be included in notifications on an opt-in basis.

– Brief statement, no mechanism or setting name given. 26.11.59

04 **Role-less users counted in org stats** Not checked yet IMPROVED 23

Statistics page organization totals now include users without a role.

– Simple UI change with no further detail. 26.11.59

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback

Connected repositories now support Composer, Hugging Face, RPM, and Swift (2026-11-23)

BLOG ↗

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

NEW INTEGRATION Connected repositories: Composer support

WAS THIS USEFUL?

StepSecurity Feedback ⓘ

Secure Registry: Automatic Cooldown Exemptions for Security Fixes, and Typosquatting Pro

BLOG ↗ RANK

StepSecurity is a security platform that hardens GitHub Actions CI/CD pipelines by detecting and preventing supply chain attacks and unauthorized network access.

> StepSecurity's Secure Registry gained automatic cooldown exemptions for security fixes sourced from multiple vulnerability databases, plus a new beta typosquatting protection feature for PyPI and Maven packages.

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

01 **Automatic cooldown exemptions for security fixes** Not checked yet

NEW 75

Advisories sourced from GitHub reviewed advisories, RustSec, and the Go vulnerability database now qualify releases for automatic exemption from the Cooldown Period. New Exempted and History views show which security fixes are currently exempted and why each advisory from the last 30 days did or did not qualify, and manual and automatic exemptions are now managed together in the Cooldown Exemptions panel, opened via Configure under Cooldown Period.

Enable automatic cooldown exemptions for security fixes so Dependabot or Renovate PRs are not held back by your cooldown period.

> 📍 In the Secure Registry console, go to Cooldown Period › Configure › Cooldown Exemptions, set the minimum severity threshold, and save. Releases fixing vulnerabilities at or above that severity will skip the Cooldown Period automatically.

Review why a recent advisory did or did not exempt a release from cooldown, covering the last 30 days.

> 📍 In the Secure Registry console, open the Cooldown Exemptions panel via Configure under Cooldown Period, then select the History view to see per-advisory qualification details.

— Names exact UI path and data sources but no severity defaults given

Secure Registry: Automatic Cooldown Exe…

02 **Typosquatting protection for PyPI and Maven** Not checked yet NEW

60

Typosquatting Protection for PyPI and Maven BETA blocks packages whose names closely resemble popular packages, with an exact-name whitelist to handle false positives.

Block typosquatted PyPI or Maven packages that closely resemble popular ones, while whitelisting known legitimate exact-name matches.

> 📍 In the Secure Registry console, enable Typosquatting Protection BETA for PyPI or Maven, then add any exact package names that triggered false positives to the exact-name whitelist.

— Clear mechanism and UI path, but beta scope and detection logic unspecified

Secure Registry: Automatic Cooldown Exe…

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Dalfox

Feedback ⓘ v2.13.1 NOTES ↗ CODE ↗ RANK

Dalfox is an XSS scanner that automates parameter analysis, reflected/stored/DOM-based XSS detection, and WAF fingerprinting for web apps.

> Dalfox v2.13.1 hardens its built-in scan server with mandatory API-key enforcement, host restrictions, and safer defaults, and adds domain-based scope filtering for scans.

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

01 Server mode security hardening

Not checked yet BREAKING

93

`dalfox server` now binds to `127.0.0.1` by default instead of `0.0.0.0`, with a startup warning when bound elsewhere without `--api-key`, and `--api-key` is required regardless of `--type`. Without a key, raw request/response bodies are dropped from results and `evidence` is trimmed to line numbers and the payload. The new `--allowed-hosts` flag restricts which hostnames the server accepts without an API key, rejecting mismatched `Host` values and DNS-rebound names, and `cookie-from-raw` is now ignored when supplied by API callers while `worker`, `timeout`, `delay`, and `trigger` are clamped to safe ranges in server mode. `POST /scan` now also requires the `Content-Type: application/json` header.

Expose the scan API on all interfaces for container or remote use while keeping it locked down with an API key.

```
$ dalfox server --host 0.0.0.0 --api-key mysecretkey
```

Submit a scan job to a locally running server instance that now requires a JSON content type header.

```
$ curl -X POST http://127.0.0.1:6664/scan -H 'Content-Type:
application/json' -H 'X-API-Key: mysecretkey' -d '{"url":
"https://example.com/search?q=test"}'
```

— Names every flag, default, and endpoint requirement changed v2.13.1

02 Out-of-scope domain filtering for scans

Not checked yet NEW

75

Scan targets can now be filtered by domain using `--out-of-scope` (a list of domains) or `--out-of-scope-file` (a file of domains to exclude), applied across URL, file, pipe, HAR, and raw-HTTP scan modes.

Scan a list of URLs while excluding third-party or out-of-scope domains so findings stay scoped to your own assets.

```
$ dalfox scan urls.txt --out-of-scope-file excluded-domains.txt --limit-result 50 --limit-result-type v
```

— Exact flags and scan modes named, with example command v2.13.1

└──▷ **BREAKING ON UPGRADE**

! To reach `dalfox server` from another machine or from outside a container, explicitly pass `--host 0.0.0.0` and set `--api-key`; the server now binds to `127.0.0.1` by default.

! Without an API key, clients that address the server by a hostname other than `localhost` or the value of `--host` must now supply `--allowed-hosts` listing those names.

! `POST /scan` now requires the `Content-Type: application/json` header; requests that omit it will be rejected.

WAS THIS USEFUL? 👍 👎

Wapiti

Feedback ⓘ 3.3.5 NOTES ↗ RANK

Wapiti is a web vulnerability scanner written in Python3

> Wapiti 3.3.5 improves its headless crawler to reach JavaScript-only content by interacting with forms and modals.

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

01 **Headless crawler fills forms and opens modals** Not checked yet

IMPROVED 35

Wapiti's headless browser crawler now fills form fields and opens modals in headless mode to reach content that only appears through JavaScript interaction, improving coverage of JS-only forms.

— Describes behaviour change but no flags, commands or scope details. 3.3.5

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ Rust SDK 0.13.3 (2026-11-23) BLOG ↗ RANK

Fastly is a cloud edge platform that accelerates content delivery, provides DDoS protection, and enables real-time CDN, security, and streaming services globally.

> Fastly's Rust SDK 0.13.3 adds the ability to read an HTTP response's reason phrase, either as raw bytes or as a lossy UTF-8 string.

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

01 **Read HTTP reason phrase in Rust SDK** Not checked yet NEW

86

The Rust SDK now supports reading a response's HTTP reason phrase via `Response::get_reason_phrase` and `ResponseHandle::get_reason_phrase`, which return a `Vec<u8>` that preserves obs-text and is empty when no reason phrase exists. A lossy UTF-8 `String` variant is also available via `get_reason_phrase_lossy`.

— Names exact methods and return types, but no broader usage context

Rust SDK 0.13.3 (2026-11-23)

WAS THIS USEFUL? 👍 👎

◆ Mobile App Security

jadx

Feedback ⓘ CVE-2026-114938 CVE ↗

Jadx is a dex to Java decompiler

`CVE-2026-114938`: Java4Shell (CVE-2026-114938) (CVSS 8.1, affects `all releases with bundled jre, released before 2026-11-18`) Not checked yet

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ

Elcomsoft Quick Triage 2.3 adds disk image mounting, direct archive browsing, and smart

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 Quick Triage 2.3 adds native handling of forensic disk images and archives alongside new triage organization and an updated interface.

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

01 **Disk image and VM mounting in Quick Triage** Not checked yet NEW

45

Quick Triage 2.3 adds native ingestion and mounting for forensic disk images and virtual machine containers, allowing direct access without separate mounting tools.

– Describes capability but no supported formats or commands named

Elcomsoft Quick Triage 2.3 adds disk im…

02 **Direct ZIP and TAR archive browsing** Not checked yet NEW

40

Quick Triage 2.3 adds direct browsing of `ZIP` and `TAR` archives without requiring decompression first.

– Names archive formats but no further mechanism detail

Elcomsoft Quick Triage 2.3 adds disk im…

03 **Smart folders with four triage categories** Not checked yet NEW

30

Quick Triage 2.3 introduces smart folders organized into four triage categories to help structure evidence review.

– Gives category count but not what the categories are

Elcomsoft Quick Triage 2.3 adds disk im…

04 **Reworked project tree and task interface** Not checked yet IMPROVED

20

Quick Triage 2.3 reworks the project tree and active tasks interface.

– Bare mention of UI rework with no specifics Elcomsoft Quick Triage 2.3 adds disk im…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.16.13 NOTES ↗ CODE ↗

Netcap is a framework for secure and scalable network traffic analysis - https://netcap.io

Netcap adds parallel small-upload scheduling via `NETCAP_AUXILIARY_CAPTURE`, MCP retry signalling with `Retry-After` and `retry_after_ms`, six-language UI localization, and a macOS fsync performance improvement cutting large capture times from 835 s to under a minute.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.16.13 https://github.com/dreadl0ck/netcap.git
# already have the repo? check out this version:
$ git checkout v0.16.13
```

Control parallel small-upload scheduling alongside large captures with the `NETCAP_AUXILIARY_CAPTURE` environment variable — set it to `0` for serial dispatch (default allows up to eight concurrent small uploads per large capture).
Not checked yet

Netcap's MCP server now signals busy responses with a `Retry-After` header and appends `; retry_after_ms=N` to tool errors so MCP clients can back off automatically. Not checked yet

The rules engine now accepts dotted paths such as `Questions.Name` in `group_by` thresholds, and those thresholds now work without `capture_time` and in bounded state. Not checked yet

Alert views now stream stores of any size through the paged API; only the 1,048,576-row and 1 MiB-per-alert limits return 413, and each 413 names the specific limit hit.
Not checked yet

Timeline metadata degradation is now reported as `metadataTruncated` and `eventsMissingMetadata` fields, with filtered responses listing `limitations`, instead of failing the session. Not checked yet

+4 more

UPDATED DETECTION SYN Scan Detection, Multiple External Connections, DNS Beaconing, DNS Tunneling threshold logic

UPDATED DETECTION SNMP community string detection reclassified

UPDATED SUPPORTED PLATFORM Audit file format support for VRRP, ApplePktap, GTPv1U/GTP

WAS THIS USEFUL? 👍 👎

Ettercap

Feedback ⓘ CVE-2026-131855 CVE ↗ STABLE

Ettercap Project

`CVE-2026-131855`: Heap out-of-bounds read in the VNC dissector when auth-type and challenge share one packet (CVSS 4.3, affects `<= 0.8.5.1`) Not checked yet

▾ **1 more release in this issue** · 2026-11-15

> CVE-2026-131854 CVE ↗ STABLE
>
> `CVE-2026-131854`: Heap out-of-bounds read in the MySQL dissector's server seed and login credential parsing (CVSS 6.5, affects `<= 0.8.5.1`) Not checked yet

WAS THIS USEFUL? 👍 👎

kubeshark

Feedback ⓘ CVE-2026-129235 CVE ↗ STABLE

Kubeshark is an eBPF-powered network observability for Kubernetes. Indexes L4/L7 traffic with full K8s context, decrypts TLS without keys.

`CVE-2026-129235`: MCP `download_file` writes fetched bytes to an unvalidated destination path (CVSS 5.3, affects `>= 53.1.0, < 53.4.0`) Not checked yet

▾ **1 more release in this issue** · 2026-08-13

> CVE-2026-129236 CVE ↗ STABLE
>
> `CVE-2026-129236`: MCP `start_kubeshark` argument injection allows redirecting deployed images to an attacker-controlled registry (CVSS 7.2, affects `>= 53.1.0, < 53.4.0`) Not checked yet

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ New Azure Storage checks: NOTES ↗ RANK

Plerion is a cloud security platform that provides risk-based prioritization and remediation guidance across cloud infrastructure, helping teams identify and address critical vulnerabilities efficiently.

> Plerion shipped four new Azure Storage posture checks, split asset page findings into configuration and event tabs, overhauled its S3 public-access check, and added browser-based MCP server sign-in from Kiro CLI.

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

01 Browser sign-in to MCP server from Kiro CLI

Not checked yet NEW

85

Connect to the Plerion MCP server from Kiro CLI using your Plerion account with no API key needed: add the Plerion MCP server URL to Kiro, run `/mcp`, sign in via the browser that opens, and select a tenant.

Authenticate Kiro CLI to the Plerion MCP server without an API key — useful for teams where individual Plerion accounts should govern access rather than a shared token

📍 In Kiro, add the Plerion MCP server URL, then run `/mcp`, sign in to Plerion in the browser that opens, and choose your tenant.

— Exact command and sign-in steps given Sign in to the Plerion MCP server from …

02 S3 public-access evaluation overhaul in PLERION-AWS-140

Not checked yet BREAKING

80

Check `PLERION-AWS-140` now evaluates every `Deny` in the bucket policy and applies Block Public Access settings the way S3 itself does, failing a bucket only when anyone or any AWS account can still list, read, or write it. This is a breaking change: some buckets will flip pass/fail result on the next scan.

— Named check with mechanism and migration impact New Azure Storage checks:

03 New Azure Storage posture checks

Not checked yet NEW

65

Four new checks were added: `PLERION-AZURE-521` flags Azure Storage accounts with no delete lock; `PLERION-AZURE-522` flags accounts that replicate blobs to an account outside their subscription; `PLERION-AZURE-523` flags accounts granting access to an external identity; and `PLERION-AZURE-524` flags Azure activity log storage with no immutability policy.

— Names all four checks but no deeper mechanism New Azure Storage checks:

---

04 **Configuration and Event findings tabs on asset pages**

Not checked yet IMPROVED 65

Asset pages now split findings into two tabs — Configuration findings (posture checks) and Event findings (CDR findings, newest first) — and the Source filter now also narrows the Event findings tab.

— Describes new UI structure and filter behavior New Azure Storage checks:

THINNER COVERAGE BELOW

05 **Retirement of Glacier vault logging check** Not checked yet

DEPRECATED 55

Check `PLERION-AWS-842` (Ensure Amazon S3 Glacier vaults have access logging enabled) is retired; its passed findings will disappear on the next scan.

— Names the retired check and its effect New Azure Storage checks:

└──▷ **BREAKING ON UPGRADE**

! Check `PLERION-AWS-842` (Ensure Amazon S3 Glacier vaults have access logging enabled) is retired; its passed findings will disappear on the next scan.

! Check `PLERION-AWS-140` evaluation logic has changed — some buckets will change result (pass or fail) on the next scan.

WAS THIS USEFUL? 👍 👎

Qualys TotalCloud

Feedback ⓘ changelog · Nov 23, 2026

BLOG ↗ RANK

Qualys TotalCloud is a cloud-native security platform providing continuous visibility, risk assessment, and compliance monitoring across multi-cloud environments and workloads.

> Qualys TotalCloud introduces TotalAI, a new module for inventorying AI software, workloads, and containers across cloud and Kubernetes environments, adding dedicated inventory views, CI/CD-to-deployment correlation, and TruRisk-based prioritization for AI supply chain risk.

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

01 **Code Assets CI/CD-to-deployment correlation** Not checked yet NEW

75

Code Assets link CI/CD repositories to built images when QScanner runs in a CI/CD pipeline, then correlates those images with container, cloud, and Kubernetes inventory to show where the software is deployed.

Trace a vulnerable AI library from its source repository all the way to the running containers in Kubernetes to understand blast radius.

📍 In TotalAI, open Code Assets, select the repository, and follow the linked image to the Containers and Kubernetes inventory view to see every deployment of that image.

— Names QScanner and traces repo-to-deployment path clearly 2026-11-23-59981c74

02 **AI Workloads inventory view** Not checked yet NEW

65

A new AI Workloads view in TotalAI shows whether AI software is installed, running, or communicating on each host, surfacing the binary, arguments, and last-seen time for each detection, and supports Qualys Query Language (QQL) search to filter and scope hosts.

Identify all hosts where AI software is actively communicating, so you can scope exposure after a new AI supply chain advisory.

📍 In TotalAI, open the AI Workloads view and use QQL to filter hosts by communication state, then review the binary and arguments columns to confirm scope.

— Names concrete fields and QQL search, but no limits given 2026-11-23-59981c74

THINNER COVERAGE BELOW

03 **AI Containers inventory** Not checked yet NEW 41

AI Containers adds installed software, Kubernetes metadata, and linked hosts and containers to give visibility inside AI container workloads.

— Describes what it shows but no UI path or mechanism 2026-11-23-59981c74

04 **TruRisk-based AI risk prioritization** Not checked yet NEW 20

TotalAI combines TruRisk scores with workload context to help prioritize AI supply chain response.

— Thin prose, no mechanism or concrete scoring detail given 2026-11-23-59981c74

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.52.0 NOTES ↗ CODE ↗ RANK

This repo hosts Prowler OSS (CLI and Local Server). Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment for Humans and AI Agents.

> Prowler 5.52.0 adds scanning for Microsoft 365 tenants in sovereign and air-gapped clouds, six new security checks across AWS, Azure, and Google Cloud, broader IPv6 exposure detection, and more surgical auto-remediation fixers.

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

01 **M365 government, sovereign, and custom cloud scanning**

Not checked yet NEW 85

A new `M365Custom` region lets Prowler scan Microsoft 365 tenants in private, sovereign, or air-gapped clouds by supplying `--m365-authority-host` and `--m365-graph-endpoint` to target custom Microsoft Graph endpoints.

Scan a Microsoft 365 tenant in a private or air-gapped cloud by supplying your own authority host and Graph endpoint.

```
$ prowler m365 --region M365Custom --m365-authority-host
https://login.microsoftonline.us --m365-graph-endpoint
https://graph.microsoft.us
```

— Runnable command with exact flags given 5.52.0

02 **Six new checks across AWS, Azure, and Google Cloud**

Not checked yet NEW 80

New checks added: `backup_vaults_lock_enabled` (verifies AWS Backup vaults have Vault Lock enabled to prevent deletion of recovery points), `sagemaker_full_access_policy_attached` (verifies no IAM role has the AWS managed `AmazonSageMakerFullAccess` policy attached, excluding service-linked roles), `aks_cluster_node_os_upgrade_channel_enabled` (verifies AKS clusters use an AKS-managed node OS auto-upgrade channel, `SecurityPatch` or `NodeImage`), `aks_cluster_kms_cmk_encryption_enabled` (verifies AKS clusters encrypt Kubernetes secrets in etcd with a customer-managed Azure Key Vault key), `databricks_workspace_uses_private_endpoint` (verifies Azure Databricks workspaces have an approved private endpoint connection), and `cloudsql_instance_deletion_protection_enabled` (verifies Google Cloud SQL instances and read replicas have deletion protection enabled).

— Each check named and explained but no run command 5.52.0

03 **More surgical remediation fixers** Not checked yet IMPROVED

75

Fixers are now more precise: EC2 security group fixers revoke only the `0.0.0.0/0` and `::/0` sources of the offending rule while keeping private CIDRs, prefix lists, and security group references intact; the `s3_bucket_policy_public_write_access` fixer removes only the statements granting public write access instead of deleting the entire bucket policy; and the `sqs_queues_not_publicly_accessible` fixer keeps non-public statements of the queue policy instead of replacing the whole policy.

— Before/after behavior named per fixer, no new command shown 5.52.0

04 **IPv6 unrestricted exposure detection across cloud checks**

Not checked yet IMPROVED

60

IPv6 internet exposure (`::/0`) is now reported as unrestricted in Google Cloud SSH/RDP firewall checks, Oracle Cloud SSH/RDP security list and NSG checks, Alibaba Cloud ECS SSH/RDP checks, and the `eks_cluster_not_publicly_accessible` check for public EKS endpoints.

— Names affected checks but no user action required 5.52.0

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - November 23, 2026

BLOG ↗ RANK

Sysdig Secure is a cloud-native security platform that provides runtime threat detection, vulnerability management, and compliance monitoring for containers, Kubernetes, and cloud environments.

> Sysdig Secure's CLI Scanner gains SARIF 2.1.0 export and IBM Cloud IaC scanning support.

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

01 **SARIF output format for CLI Scanner** Not checked yet NEW 83

The Sysdig CLI Scanner can output scan results in SARIF 2.1.0 format, either to the console or to a file, using `sarif` or `sarif-file=<path>` values with the `--output` option. This enables upload of results to GitHub code scanning or GitLab via `artifacts:reports:sarif`.

Export a scan as SARIF to a file so it can be uploaded to GitHub code scanning as a code-scanning artifact in CI.

```
$ sysdig-cli-scanner --output sarif-file=results.sarif <target>
```

— Names exact flag, values, and CI integration targets Secure SaaS - November 23, 2026

THINNER COVERAGE BELOW

02 **IBM Cloud IaC scanning support** Not checked yet NEW 30

Sysdig CLI Scanner now supports scanning IBM Cloud resources in IaC mode.

— Bare statement of support, no detail on scope or usage Secure SaaS - November 23, 2026

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Chainguard

Feedback ⓘ Week of 2026-11-23 NOTES ↗

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

Chainguard adds the Athena Partner API so shield partners can register mitigations and report finding statuses via chainctl or REST, replacing hand-reviewed capability manifests.

NEW DETECTION Partners register mitigation methods and report `MITIGATED`, `DECLINED`, or `FAILED` statuses per finding via the new Athena Partner API, available through `chainctl` or REST. Not checked yet

Approved partners receive findings through the Athena OSV API scoped only to vulnerabilities they can mitigate, replacing hand-reviewed capability manifests. Not checked yet

NEW SUPPORTED PLATFORM New container images

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Protection for Linux

Feedback ⓘ

Version 5.23.3New NOTES ↗ RANK

Sophos Protection for Linux is a security agent that provides real-time malware detection, threat prevention, and endpoint protection for Linux systems.

> Sophos Protection for Linux adds Tamper Resiliency to detect and recover unauthorized changes to agent files, expands `sophosctl` for local recovery and log-level management, introduces the Endpoint Self Help diagnostic CLI, enables Update Cache and Message Relay roles for Linux servers, broadens Sophos Data Lake telemetry with DNS/SNI, container, cleanup and detection-context events, and adds kernel 7.1 support.

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

01 **Expanded sophosctl local recovery commands** Not checked yet

IMPROVED 85

Administrators can trigger a product update, force a reinstall, or view, set, or reset global or component-specific log levels locally using expanded `sophosctl` commands — providing a supported recovery path without manual config file edits.

— Names exact tool and functions but omits precise flag syntax Version 2026.4New

02 **Tamper Resiliency for SPL Agent files** Not checked yet IMPROVED 65

Adds automatic detection and recovery for unauthorized changes to SPL Agent files via Tamper Resiliency, which now also covers additional plugin registry files and detects unexpected files in protected installation directories.

— Explains mechanism and expanded scope but no config flag given Version 2026.4New

Version 2026.3

03 **Expanded Data Lake telemetry and detection context**

Not checked yet NEW 60

Sophos Data Lake now receives DNS activity and Server Name Indication (SNI) hostname events for expanded network visibility, container information added to supported detection events, cleanup activity events so analysts can correlate remediation activity with antivirus detections, and additional file, process, script, network, and authentication evidence in detection events.

— Lists telemetry types added but no query or access detail Version 2026.3

Version 2026.4New

THINNER COVERAGE BELOW

04 **Update Cache and Message Relay roles for Linux servers**

Not checked yet NEW 55

Linux servers with SPL Agent installed can now be configured as Update Caches and Message Relays in Sophos Central.

— Names roles and console but no step-by-step config shown Version 2026.3

05 **Endpoint Self Help diagnostic CLI** Not checked yet NEW

45

A new command-line tool, Endpoint Self Help, lets administrators view product health and collect diagnostic information from a Linux server.

— Names the tool but no command syntax given Version 2026.3

06 **Kernel 7.1 support** Not checked yet NEW 25

Adds support for kernel 7.1.

— Bare compatibility note with no further detail Version 2026.3

NEW SUPPORTED PLATFORM Extended kernel support

Version 5.23.3New NOTES ↗ LATEST STABLE

Sophos Protection for Linux 5.23.3 adds Helm-based process exclusions and configurable data-table controls, plus `/proc/<pid>/comm` rule filtering and Runtime Detection integrity checks.

Exclude selected executables from process, connection, and DNS events sent to the data lake by adding Helm configuration options for process exclusions, reducing event volume in container environments. Not checked yet

Turn on supported data tables required for deployment and investigation by configuring expanded Helm values for data-table selection. Not checked yet

> Detection rules now use process-name information from `/proc/<pid>/comm` for more precise threat filtering. Not checked yet
>
> Runtime Detection files now include integrity checks to improve resilience against unexpected modification. Not checked yet

WAS THIS USEFUL? 👍 👎

Sophos Central

Feedback ⓘ November 2, 2026 NOTES ↗ RANK

Sophos Central is a cloud-based management platform that enables administrators to manage and monitor Sophos security products, endpoints, and policies from a single console.

> Sophos Central strengthened password security in Sophos Fusion by blocking commonly used passwords.

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

01 **Common password blocking in Sophos Fusion** Not checked yet

IMPROVED 35

Sophos Fusion now disallows passwords that appear on public lists of common passwords, reducing credential-stuffing exposure.

— Explains the change but no mechanism, scope, or config detail given. November 2, 2026

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Splunk Security Content

Feedback ⓘ v6.11.1 NOTES ↗

Splunk Security Content

Adds a `zeek_files` macro and updates the `linux_shells` macro (removing `eshell`) to support new web-layer analytics.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v6.11.1
https://github.com/splunk/security_content.git
# already have the repo? check out this version:
$ git checkout v6.11.1
```

Adds a `zeek_files` macro to support new web-layer analytics. Not checked yet

Updates the `linux_shells` macro and removes the `eshell` entry. Not checked yet

NEW DETECTION HTTP Large POST and Oversized Multipart Header Activity

NEW DETECTION Linux Auditd Shell from Web or Application Directory

NEW DETECTION Web Anomaly Followed by Server Process Execution

UPDATED DETECTION Updated Analytics

NEW DETECTION New Analytic Stories

└──▷ BREAKING ON UPGRADE

! Starting in ESCU 6.12.0, Splunk Enterprise 9.0, 9.1, 9.2, 9.3 and Splunk Cloud Platform 10.1 will be removed from ESCU compatibility listings.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Lookyloo

Feedback ⓘ CVE-2026-135617 CVE ↗ RANK

Lookyloo is a web interface that allows users to capture a website page and then display a tree of domains that call each other.

> Lookyloo v1.43.1 adds a separate public monitoring URL with improved nginx setup and ships a broad redesign of the web UI, touching nearly every page from the home page to pivot views.

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

01 **Redesigned web UI across site and capture pages** Not checked yet

IMPROVED 60

Lookyloo's web UI is redesigned: a new site header, footer, page headers, and home page; restyled capture list, category tags, table controls, loader, and credits page; restyled capture form and search, with a tutorial added to every page except the tree; blurred screenshots shown in capture lists; a restyled statistics page; custom 404 and 500 error pages and a View capture button; and redesigned pivot pages, with the Explore and Statistics sections disabled pending further work.

— Names every reworked UI area but gives only navigation, no commands v1.43.1

THINNER COVERAGE BELOW

02 **Separate public URL for monitoring** Not checked yet NEW

37

Lookyloo now exposes a separate public URL dedicated to monitoring and improves the nginx setup that sits behind it.

— Names nginx but no exact config paths or endpoints v1.43.1

> CVE-2026-135617 CVE ↗ STABLE
>
> `CVE-2026-135617`: Lookyloo Inconsistent Authorization Enforcement for Private Captures Allows Unauthorized Data Access (affects `<= 1.43.0`) Not checked yet

WAS THIS USEFUL? 👍 👎

MISP

Feedback ⓘ v2.5.53 NOTES ↗

MISP (core software) - Open Source Threat Intelligence and Sharing Platform

MISP adds opt-in per-user module credentials controlled by `Plugin.Modules_allow_personal_credentials`, with secrets stored in `module_api_keys` and rotated via `cake Admin reencrypt`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.5.53 https://github.com/MISP/MISP.git
# already have the repo? check out this version:
$ git checkout v2.5.53
```

Enable per-user API keys, account names, and passwords for enrichment, import, and export modules by turning on `Plugin.Modules_allow_personal_credentials` (off by default); secrets are stored in the `module_api_keys` user setting, encrypted individually, masked in the UI and audit log, and rotated with `cake Admin reencrypt`. Not checked yet

The REST API client now documents the warninglist add, edit, delete, export, and import endpoints. Not checked yet

NEW INTEGRATION OffSeq Mirage community feed

UPDATED DETECTION Galaxies, objects, taxonomies and warninglists updated to latest versions

WAS THIS USEFUL? 👍 👎

Watcher

Feedback ⓘ v3.8.0 NOTES ↗ RANK

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

> Watcher v3.8.0 replaces its local summarization models with a multi-provider LLM gateway (OpenAI, Anthropic, Gemini, Ollama, or any LiteLLM-compatible endpoint), adds on-demand regeneration of AI summaries with a configurable cooldown, richer threat extraction in its reports, new notification layouts, and small UI preference additions.

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

01 **Multi-LLM gateway via LiteLLM** Not checked yet NEW 100

New `llm_gateway` app built on LiteLLM supports OpenAI, Anthropic, Gemini, Ollama, and any LiteLLM-compatible endpoint as the AI backend, configured via `deployment/.env` keys `OPENAI_API_KEY`, `OPENAI_MODEL`, `OPENAI_BASE_URL`, `ANTHROPIC_*`, `GEMINI_*`, `OLLAMA_API_KEY`, `OLLAMA_MODEL`, `OLLAMA_API_URL`, `GENERIC_LLM_URL`, `GENERIC_LLM_KEY`, `GENERIC_LLM_MODEL` (or directly from the Connectors page). It falls back automatically to a local Hugging Face model (override via `HF_LOCAL_MODEL`, default `Qwen/Qwen2.5-0.5B-Instruct`) when no Default LLM is selected, the API is unreachable, or no key is configured, with TLS verification controlled via `LITELLM_VERIFY_SSL` and `SSL_CERT_FILE`; a connector is set as the Default LLM via `POST /api/connectors/<id>/set-default-llm/` (superusers only).

Point Watcher at an internal Ollama instance as the LLM backend for AI summaries, bypassing a corporate proxy.

```ini
OLLAMA_MODEL=llama3
OLLAMA_API_URL=http://ollama.internal:11434
NO_PROXY=ollama.internal
LITELLM_LOCAL_MODEL_COST_MAP=True
```

— Full env-var list, fallback mechanism and endpoint all named v3.8.0

02 **On-demand summary regeneration with cooldown** Not checked yet NEW 90

Adds `POST` endpoints `/api/threats_watcher/summary/regenerate-weekly/`, `/api/threats_watcher/summary/by-keyword/<keyword>/regenerate/`, and `/api/threats_watcher/summary/by-keyword/<keyword>/regenerate-breaking/` to trigger immediate regeneration of the weekly report, keyword summaries, and breaking news. A global cooldown (1–1440 minutes) is

readable/updatable via `/api/threats_watcher/summary/config/` (`GET` for all authenticated users, `POST` for superusers), surfaced in the UI as a Regenerate button with a superuser-configurable cooldown slider (1 minute to 24 hours, default 15 minutes) and a remaining-time indicator on the weekly report, breaking news and keyword summary views.

Trigger an immediate regeneration of the weekly threat report from the API, e.g. from a CI pipeline or a scheduled job.

```
$ curl -s -X POST
https://watcher.example.com/api/threats_watcher/summary/regenerat
e-weekly/ \
  -H 'Authorization: Token <your_token>'
```

Set the global regeneration cooldown to 30 minutes so analysts cannot trigger back-to-back LLM calls.

```
$ curl -s -X POST
https://watcher.example.com/api/threats_watcher/summary/config/ \
  -H 'Authorization: Token <superuser_token>' \
  -H 'Content-Type: application/json' \
  -d '{"value": 30}'
```

— Three endpoints and cooldown API given with working curl examples v3.8.0

---

03 **Upgrade and deployment changes** Not checked yet BREAKING

70

Database migrations must now be applied manually by running `make migrate` from `deployment/` after starting the new images, since the `watcher` container no longer applies them automatically on startup. The Watcher Default UI theme is replaced by Darkly, with affected users switched automatically. Corporate proxy users must add internal LLM hostnames to `NO_PROXY` / `no_proxy` in `deployment/.env` before running `make up`, otherwise LLM requests are routed through the proxy and may be rejected.

— Exact commands and config keys given for required upgrade steps v3.8.0

---

04 **Richer extraction in breaking news, keyword summaries and weekly report**

Not checked yet IMPROVED

65

Breaking news now extracts impact level, targeted sectors, threat actors, malware and CVEs with an alert score; keyword summaries extract CVEs, APT groups, malware,

MITRE ATT&CK techniques, sectors and countries with a relevance score based on the organisation; the weekly report groups top threats into at most 5 severity clusters with an executive summary and the first publication date of each threat.

— Describes new fields extracted but no API or config to act on v3.8.0

05 **Local summarization models removed for LLM gateway**

Not checked yet DEPRECATED 60

The local BART and FLAN-T5 summarization models are removed; a Default LLM connector must now be configured on the Connectors page, or every summary falls back to the local Hugging Face model (`Qwen/Qwen2.5-0.5B-Instruct` by default).

— Clear migration requirement named, no further mechanism detail v3.8.0

THINNER COVERAGE BELOW

06 **Configurable weekly report schedule** Not checked yet BREAKING 55

The weekly report schedule is now configurable via `WEEKLY_SUMMARY_DAY` and `WEEKLY_SUMMARY_HOUR` in `deployment/.env`, defaulting to Thursday at 09:30, changing from the previous Monday schedule.

— Env vars and default named, mechanism beyond that is thin v3.8.0

07 **New notification layouts for reports** Not checked yet NEW 50

Adds new notification layouts for the weekly report and breaking news: Slack attachments with severity colours, email-safe HTML, and Matrix HTML for Citadel.

— Names three output formats but no config or endpoint given v3.8.0

08 **Interface zoom and typewriter effect for summaries**

Not checked yet NEW 50

Adds interface zoom levels (75%, 80%, 90%, 100%) saved to the user profile, and a typewriter-effect toggle for AI summaries.

— Exact zoom percentages given but a thin UI preference v3.8.0

└──▷ BREAKING ON UPGRADE

! Database migrations are required: run `make migrate` from `deployment/` after starting the new images — the `watcher` container does not apply them automatically on startup.

! The `Watcher Default` UI theme is replaced by `Darkly`; users who had it selected are switched automatically.

! The weekly report is now generated on Thursday at 09:30 instead of Monday.

! The local BART and FLAN-T5 summarization models are removed; a Default LLM connector must be configured on the Connectors page, or every summary falls back to the local Hugging Face model (`Qwen/Qwen2.5-0.5B-Instruct` by default).

! Corporate proxy users must add internal LLM hostnames to `NO_PROXY` / `no_proxy` in `deployment/.env` before running `make up`, otherwise LLM requests are routed through the corporate proxy and may be rejected.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Elastic Security Labs

Feedback ⓘ

Press play on Elastic Security 9.6, banger after banger (2026-11-23) BLOG ↗

RANK

Elastic Security Labs is a powered by Elasticsearch. Ready to deploy your data?

> Elastic Security introduced AlertZero, a technical-preview AI SOC system with six autonomous Workers spanning alert triage, threat hunting, detection engineering and forensics, complete with tunable autonomy levels, per-step model tiers, cross-Worker handoffs, a proposed-actions review page and case escalation.

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

01 **Proposed actions review page** Not checked yet NEW 85

Analysts review and approve proposed actions from a dedicated proposed actions page grouped under Respond, Investigate, and Configure, with each action backed by the Worker's gathered evidence and reasoning inside an Elastic Agent Builder investigation conversation.

Walk through approving or declining an AlertZero proposed action — for example, a host isolation recommended by the Endpoint Analysis Worker — before it runs against your environment.

> 📍 In Elastic Security, go to AlertZero › Proposed Actions, select the action under Respond, review the evidence and reasoning in the attached investigation, then click Approve (runs as you, with your privileges) or Decline.

— Exact navigation and approve/decline steps provided

AlertZero in Elastic Security: Every al…

02 **AlertZero AI Workers for SOC automation** Not checked yet

NEW DETECTION NEW SUPPORTED PLATFORM NEW INTEGRATION NEW

73

AlertZero ships in technical preview on Elastic Cloud Serverless with Elastic Security 9.6 (ECH and Self-Managed support expected in early December), bringing six AI Workers across four Watches to automate alert triage, attack discovery, continuous threat hunting, rule tuning, rule coverage gap review, and endpoint forensic analysis.

— Names product, version, workers and watches but no enable steps

AlertZero in Elastic Security: Every al…

03 **Autonomy levels for AlertZero Workers** Not checked yet NEW

65

Each Worker's autonomy can be set to Manual, Assisted, or Supervised to control whether proposed actions — such as closing false positives, isolating hosts, or editing detection rules — wait for analyst approval or run automatically.

— Mechanism and examples given, no exact config path

AlertZero in Elastic Security: Every al…

04 **Alert Triage Worker confidence scoring** Not checked yet NEW

65

The Alert Triage Worker classifies new alerts as true or false positives with a confidence score, and at Supervised autonomy auto-closes false positives at or above an 85% default confidence threshold.

— Concrete 85% threshold named, no tuning command given

AlertZero in Elastic Security: Every al…

THINNER COVERAGE BELOW

05 **Model tier assignment per Worker step** Not checked yet NEW

55

AI steps in each Worker workflow can be assigned to one of three model tiers — Efficient, Balanced, or Frontier — so the right model is used per step and can be swapped without changing the Worker's logic.

— Tiers named but no UI path to assign them AlertZero in Elastic Security: Every al…

06 **Cross-Worker handoffs** Not checked yet NEW

55

Workers hand off work to each other: Alert Triage feeds false-positive counts to the Rule Tuning Worker, Attack Discovery Worker escalates to the Endpoint Analysis Worker, and

the Continuous Threat Hunt Worker records detection gaps for the Rule Coverage Worker.

.

— Specific data flows named but purely descriptive

AlertZero in Elastic Security: Every al…

07 **Case escalation for related investigations** Not checked yet NEW

35

Related investigations can be escalated into a shared case conversation, assigned to team members, and marked private or public.

— Described in one sentence with no mechanism detail

AlertZero in Elastic Security: Every al…

NEW DETECTION MITRE ATLAS mappings for AI-platform prebuilt detection rules

NEW DETECTION New detection rules for AI agents and large language model platforms

NEW SUPPORTED PLATFORM Ransomware protection extended to Linux

NEW INTEGRATION Automatic Migration support covers Splunk, QRadar, and Microsoft Sentinel

NEW INTEGRATION Entity resolution expanded to cover CrowdStrike SIDs, Entra GUIDs, and Microsoft 365 UPNs

NEW COMPLIANCE FRAMEWORK MITRE ATLAS mappings for all 143 AI SIEM rules

NEW INTEGRATION Fleet integrations for AI platform log collection

NEW DETECTION Instruction override (prompt injection) detection model

NEW DETECTION String encoding/hash detection model

NEW DETECTION Word obfuscation detection model

NEW DETECTION Refusal detection model

NEW DETECTION Prebuilt detection rules mapped to MITRE ATLAS and OWASP Top 10 for Agentic Applications 2026

NEW DETECTION Anomaly detection jobs for AWS Bedrock — tool-request volume, rare tool selection, high call volume

NEW INTEGRATION Security for AI package integration with AWS Bedrock

NEW COMPLIANCE FRAMEWORK Detection rules mapped to OWASP Top 10 for Agentic Applications 2026

> Press play on Elastic Security 9.6, banger after banger (2026-11-23) BLOG ↗
>
> LATEST STABLE
>
> Elastic Security 9.6 adds AlertZero agentic SOC automation with four Watch types, custom YARA memory scanning, Entity Analytics AI brief, and agentic SIEM rule migration editing.

Elastic Security adds AlertZero, an agentic layer with four Watches — Alert Triage, Threat Hunt, Detection Engineering, and Forensics — each running Workers in the background that classify alerts, hunt for threats, tune noisy rules, close detection gaps, and reconstruct attack timelines, with Manual, Assisted, or Supervised autonomy levels per Worker. Not checked yet

Elastic Security adds the YARA signatures tab on the Artifacts page, letting analysts apply custom YARA signatures globally or to specific Defend policies and enabling scanning via the Memory threat protection card — a match surfaces a Memory Threat – Detected alert with the matched signature name and memory strings. Not checked yet

Elastic Security's Automatic Migration now supports agentic editing of translated detection rules from chat, covering ES|QL logic changes, custom index patterns, custom field mappings, syntax fixes, prebuilt-rule matching, and alternative rule recommendations across Splunk, QRadar, and Microsoft Sentinel. Not checked yet

Entity Analytics adds an AI brief summarising entity risk posture, naming which entities need attention, listing up to three entities to investigate first, and supporting follow-up questions in Agent Builder without rebuilding investigation context. Not checked yet

Entity Analytics adds a refreshed overview with six attention tiles — Severely alerting, Newly high-risk, Risk movers, New & alerting, Anomalous, and Watchlisted & alerting — filterable by entity type, risk level, criticality, source, or watchlist across 24-hour, 7-day, and 30-day periods. Not checked yet

+1 more

Elastic ships 143 detection rules for AI agents and LLM platforms (2026-11-23)

BLOG ↗ STABLE

Elastic 9.6 adds Elastic Defend on-host prevention for AI coding agents (Claude Code, Cursor, Codex) and a new MITRE ATLAS mapping layer across all AI SIEM rules.

Elastic Defend behavior rules now run on-host in prevent mode for coding agents, killing offending processes before actions complete — including agent-driven persistence and malicious Cursor workspace code execution. Not checked yet

All 143 AI SIEM rules ship with MITRE ATLAS technique mappings alongside MITRE ATT&CK, with Elastic 9.6 being the first release to carry those mappings into the product. Not checked yet

Thinking fast about prompt injection: New System 1 models in Elastic (2026-11-23)

BLOG ↗ STABLE

> Elastic Security adds a Security for AI Kibana integration package with four ingest-time NLP classifiers writing scores to `security_ai` fields, LLM-as-judge escalation for high-scoring events, and ECS `gen_ai` field compatibility.
>
> Four lightweight NLP classifiers score GenAI traffic at ingest time inside your own cluster, writing to `security_ai.encoding`, `security_ai.hash_string`, `security_ai.word_obfuscation`, `security_ai.instruction_override`, and `security_ai.refusal` fields with no GPU required. Not checked yet
>
> LLM-as-judge escalation is applied only to high-scoring events, keeping cost proportional to actual risk. Not checked yet
>
> Models use ECS `gen_ai.input.messages` and `gen_ai.output.messages` fields, making the package compatible with any integration that populates those fields. Not checked yet
>
> Available as a technical preview in Elastic Stack 9.6 on a Platinum or Trial subscription. Not checked yet

---

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Impact analysis in Unity Catalog lineage (Beta) (2026-11-23) BLOG ↗ RANK

Databricks is a unified data analytics platform that enables teams to collaborate on big data processing, machine learning, and AI using Apache Spark.

> Databricks shipped beta impact analysis for Unity Catalog lineage with CSV export, an upcoming ai_forecast version 2 with holidays and external covariates, GA OAuth authentication for webhook notification destinations, larger dashboard subscription attachments, and several Genie One updates including GA voice input and cross-folder skills.

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

01 Impact analysis in Unity Catalog lineage

Not checked yet NEW

75

A new Beta Impact analysis view under the Lineage tab for a table or column shows downstream assets that depend on it, each with an impact level and owner, and lets Genie Code summarize the impact and export a CSV of owners.

Find all downstream assets affected by a table change and export the owner list as a CSV for stakeholder notification.

> 📍 In the Unity Catalog UI, open the target table or column, navigate to the Lineage tab, select Impact analysis, review downstream assets and their impact levels, then use Genie Code to summarize and export a CSV of owners.

– Clear UI path and worked example, but Beta mechanism details are limited

Impact analysis in Unity Catalog lineag…

02 ai_forecast version 2 with holidays and covariates

Not checked yet NEW

75

`ai_forecast` version 2 adds holidays, external covariates, and non-negative forecast output to foundation-model-based time series forecasts in SQL; it becomes the default for workspaces with the compliance security profile enabled starting January 2027, while queries without a specified version keep using `ai_forecast` version 1 until then.

– Names the function and rollout date but no runnable example given

ai_forecast version 2 will soon be avai…

03 Voice input, push notifications, and cross-folder skills in Genie One

Not checked yet NEW 65

Genie One's voice input is now generally available on the web, desktop app, and mobile app for US and EU regions, letting users dictate prompts that are transcribed to text; the Genie mobile app now sends Android push notifications for long-running prompts and scheduled tasks; and Genie One now lists workspace skills and skills from other folders registered in Genie Code.

— Three concrete capabilities described but no setup steps given

Genie One enhancements (2026-11-19)

THINNER COVERAGE BELOW

04 **OAuth authentication for webhook destinations** Not checked yet

NEW 55

Webhook notification destinations can now authenticate using OAuth 2.0 client credentials to obtain an access token from a configured OAuth token endpoint; this is generally available.

— Describes the auth mechanism but no config key or endpoint named

OAuth authentication for webhook notifi…

05 **Dashboard subscription attachments and Genie Code styling**

Not checked yet IMPROVED 55

Dashboard subscription emails now support up to 24 MB of combined attachments, up from 9 MB, and Genie Code for dashboard authoring now reproduces the visual styling of a screenshot more closely when recreating a dashboard from it.

— Gives concrete size limits but no config path to change behavior

Dashboard enhancements (2026-11-19)

NEW INTEGRATION SharePoint Online indexing for Genie Search

NEW DETECTION ai_mask v2.0 Beta — system-defined classifiers with Databricks-built detection model

└──▷ ALSO FROM THESE RELEASES

> Search SharePoint content with Genie Search (Beta) (2026-11-23) BLOG ↗
>
> STABLE
>
> Genie Search adds SharePoint Online as a connected source BETA, with results respecting users' existing SharePoint access permissions.

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ Data Pipelines (Manage) (2026-11-23)

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 introduced Data Pipelines, a visual graph-based system for routing, parsing, enriching, and masking logs before they're indexed, with built-in testing and monitoring.

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

01 **Data Pipelines for log routing, enrichment, and masking**

Not checked yet NEW 74

A new Manage Data > Data Pipelines module lets you route logs into a pipeline using a routing expression and process them through ordered groups of processors and routers arranged in a visual graph; each log message passes through only the first matching pipeline, so pipeline order controls overlapping rules. You can test a whole pipeline or specific processing groups against sample data before publishing, and once live, monitor each pipeline's input and output data volume and health events. To build one: go to Manage Data > Data Pipelines, click `Add Pipeline`, enter a routing expression, add ordered processing groups, click `Test` with sample data, then `Publish`.

Walk through building, testing, and publishing a new Data Pipeline in the Sumo Logic UI.

> 📍 1. Go to Manage Data › Data Pipelines and click `Add Pipeline`. 2. Enter a routing expression to select which logs enter this pipeline. 3. Add ordered processing groups (parse, enrich, mask, or route processors). 4. Click `Test` and supply sample log data to validate the pipeline output before publishing. 5. Click `Publish` to activate the pipeline. 6. Open the pipeline's monitoring view to watch input/output data volume and health events.

— Full UI workflow and mechanism given, but no API/config surface named.

Data Pipelines (Manage) (2026-11-23)

WAS THIS USEFUL? 👍 👎

◆ AI SOC

Exaforce

Feedback ⓘ Meet ExaU, Our New Chief AI Officer NOTES ↗

RANK

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

> Exaforce launched ExaU, a platform for assembling named, role-specific AI agents — threat hunters, access reviewers, policy officers, security chiefs of staff, CISO briefing drafters — that operate inside its existing knowledge graph with human approval gates and full audit logging.

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

01 **ExaU platform for role-specific AI agents** Not checked yet NEW

79

ExaU lets teams build AI agents from templates or plain-language descriptions for roles such as threat hunter, access reviewer, policy officer, security chief of staff, or CISO briefing drafter, each with its own name, role, and persistent memory that can be reviewed and edited under 'lessons picked up on the job', team notes, or only you entries in the Memory tab. Any agent action affecting employees, systems, or communications requires an Allow/Deny approval in the Exaforce console, with every step written to the Exaforce audit log; agents can be shared across a team so one finding triggers a threat hunt, an activity audit, and employee outreach simultaneously, with results rolled into one assessment. Agents run through identities and connections already configured in Exaforce, requiring no new credentials, tokens, or API logins.

Review and correct what a threat-hunting agent has learned about your environment before it runs its next hunt.

📍 In the Exaforce console, open your threat-hunting agent, go to the Memory tab, and edit or remove any entry under 'lessons picked up on the job', team notes, or only you notes to adjust its behaviour going forward.

Gate a potentially disruptive agent action — such as notifying an employee — so it cannot proceed without explicit sign-off.

📍 In the Exaforce console, when an agent surfaces an action that affects employees, systems, or communications, select Allow or Deny in the approval prompt; the decision and every preceding step are written to the Exaforce audit log automatically.

— Names roles, memory tabs, approval flow and console steps clearly.

Meet ExaU, Our New Chief AI Officer

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

vulnerability-db

Feedback ⓘ v7.0.7 NOTES ↗ RANK

Vulnerability-db is a vulnerability database and package search for sources such as Linux, OSV, NVD, GitHub and npm. Powered by sqlite, CVE 5.2, purl, and vers.

> vulnerability-db v7.0.7 adds leaner CI-oriented database flavors and tightens error handling for metadata searches, alongside a breaking change to release-number query matching.

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

01 app-2y and app-2y-extended database flavors

Not checked yet NEW

80

Refresh the smallest complete databases with `vdb db refresh full --flavor app-2y` or `--flavor app-2y-extended`, suited for CI runs where download time matters most.

Pull the smallest complete database in a CI pipeline to minimise download time.

```
$ vdb db refresh full --flavor app-2y
```

— Runnable command with named flags given v7.0.7

02 Clearer metadata-search errors on incompatible flavors

Not checked yet IMPROVED

75

Metadata searches (`--search-text`, `--search-alias`, `--search-reference`, `--search-symbol`) now exit 1 with an explanation naming `-extended` flavor or `VDB_INCLUDE_METADATA=true` as remedies when run against `full` or `app-only` flavors that lack metadata tables.

— Names exact flags and env var remedy v7.0.7

THINNER COVERAGE BELOW

03 Release-number query matching changed for date-shaped ceilings

Not checked yet BREAKING

50

A release-number query no longer matches rows with date-shaped ceilings, including floor-less `<20230101` rows on non-Go packages.

— Describes behaviour change but no migration guidance v7.0.7

└──▷ **BREAKING ON UPGRADE**

! A release-number query no longer matches rows with date-shaped ceilings, including floorless `<20230101` rows on non-Go packages.

WAS THIS USEFUL? 👍 👎

oxo

Feedback ⓘ v2.25.5 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> oxo adds an `/idle` endpoint allowing agents to report whether they are currently processing a message.

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

01 **Agent idle status endpoint** Not checked yet NEW 65

A new `/idle` endpoint lets callers query whether an agent is currently handling a message, exposing its processing state.

– Names exact endpoint but lacks response format or usage detail. v2.25.5

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Fleet

Feedback ⓘ fleet-v4.96.0 NOTES ↗

Fleet is an open-source platform for managing and querying device fleets, providing real-time visibility into endpoints using osquery.

Fleet adds a Google conditional access integration and a new REST API endpoint to mark or unmark company-owned hosts as managed in Google.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch fleet-v4.96.0
https://github.com/fleetdm/fleet.git
# already have the repo? check out this version:
$ git checkout fleet-v4.96.0
```

NEW INTEGRATION Connect Google conditional access via `Settings > Integrations > Conditional access`, the REST API, or GitOps to let Google block sign-in from hosts Fleet does not manage. Not checked yet

Use `PUT /api/v1/fleet/hosts/:id/google_device_id` to mark a company-owned host as managed in Google, and `DELETE` to unmark it — both require Google conditional access to be connected. Not checked yet

Android hosts now enforce a passcode during setup (before the home screen) when an Android configuration profile includes `passwordPolicies` or `passwordRequirements` and no labels. Not checked yet

└──▷ BREAKING ON UPGRADE

! `PUT /api/v1/fleet/hosts/:id/google_device_id` now requires Google conditional access to be connected and the host to be company-owned; calls that previously succeeded without these conditions will now fail.

WAS THIS USEFUL? 👍 👎

ITarian

Feedback ⓘ

Release Notes: ITarian Platform Update - November 25, 2026 (2026-11-23) BLOG ↗

RANK

Streamline IT operations with ITarian's all-in-one IT management software. Automate tasks, enhance security, and boost efficiency.

> ITarian introduces hierarchical device grouping with automatic assignment rules and a new connection-error filter in the Device List.

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

01 Automatic device assignment rules by hostname or OS

Not checked yet NEW 95

Group Management now supports assignment rules that automatically move new devices into groups based on hostname match (`starts with` , `contains` , `ends with` , or `regex`) or OS type (Windows, Linux, macOS, Android, iOS), with a Test Expression panel to validate a rule before saving.

Automatically route newly enrolled Windows endpoints whose hostnames start with `WIN-DC` into a dedicated group, without manual sorting.

> 📍 In Group Management, open Assignment Rules for the target customer, add a rule with hostname condition `starts with` set to `WIN-DC` and OS type `Windows` , then use the Test Expression panel to verify before saving.

— Concrete match types, OS options, and a worked example

Release Notes: ITarian Platform Update …

02 Multi-level device groups with policy inheritance

Not checked yet NEW 65

Devices can now be organized into nested groups with no maximum depth using `Add Subgroup` from the group tree or the Group Management toolbar, and policies can be assigned per subgroup or inherited from the nearest parent group.

— UI path and behavior given, no API or flag detail

Release Notes: ITarian Platform Update …

THINNER COVERAGE BELOW

03 Device List filter for connection errors

Not checked yet NEW 50

The Device List can be filtered by `Client security connection error` to surface all devices sharing that connectivity issue in one step.

Quickly surface all endpoints with agent connectivity problems by filtering the Device List to `Client security connection error`.

> 📍 In the Device List, open the filter panel and select `Client security connection error` to scope the view to only affected devices.

— Clear filter name and steps, but minimal added mechanism

`Release Notes: ITarian Platform Update …`

WAS THIS USEFUL? 👍 👎

Iru `Feedback` ⓘ

`What's New in Iru: App Preservation for Return to Service, Custom Declarations, and PCI I`

NOTES ↗

Iru is a device-centric platform that unifies security, IT, and compliance around devices and everything they touch.

Iru adds App Preservation to retain managed apps through Return to Service erases, and Custom Declarations with Apple DDM access and inline JSON validation.

Retain managed apps on shared devices through a Return to Service erase using App Preservation, so apps survive device wipes without requiring reinstallation.

`Not checked yet`

Deploy Custom Declarations with direct access to Apple's DDM controls, now with inline JSON validation. `Not checked yet`

NEW COMPLIANCE FRAMEWORK PCI DSS readiness tracking

WAS THIS USEFUL? 👍 👎

◆ Patch Management

Patch My PC

Feedback ⓘ

Patch My PC Publisher Preview Version 3.26.11032.0 Released (2026-11-23) BLOG ↗

RANK

Simplify IT with Patch My PC's endpoint management software. Automate patching, app updates, reporting, and analytics to save time and boost security.

> Patch My PC's Publisher preview adds a public preview of a Remote User Interface and introduces repository path safety checks for content operations.

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

01 **Remote User Interface public preview** Not checked yet NEW

30

Patch My PC Publisher adds a public preview of the Remote User Interface, allowing administration of Publisher remotely rather than only from the local console.

— Named feature but no mechanism, UI path, or config detail given

Patch My PC Publisher Preview Version 3…

02 **Repository path safety checks for content operations**

Not checked yet IMPROVED

25

Publisher adds repository path safety checks for content operations, intended to prevent unsafe or incorrect repository paths during content processing.

— Mentioned only in summary, no specifics on checks or scope

Patch My PC Publisher Preview Version 3…

WAS THIS USEFUL? 👍 👎

◆ Attack Surface Management

Censys

Feedback ⓘ November 23, 2026 (2026-11-23) BLOG ↗ RANK

Censys is an internet intelligence platform that continuously scans and indexes internet-connected devices, exposing attack surfaces and helping security teams identify vulnerabilities.

> Censys expanded its fingerprinting coverage with nearly 300 new software, hardware, and OS fingerprints, plus a new ASM risk fingerprint targeting the nginx2root vulnerability class.

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

01 **287 new software and hardware fingerprints** Not checked yet NEW

50

Censys added 287 new software and hardware fingerprints to the platform, covering devices and services including Atlassian Jira, Apache NiFi, Palo Alto Networks Expedition, SolarWinds Access Rights Manager, Supermicro BMC, and Western Digital My Cloud NAS, among dozens more. It also added 11 new OS fingerprints, including SonicWall SonicOS, Synology Router Manager, Synology BeeStation OS, Juniper Junos OS, and Hillstone Networks StoneOS.

— Names many products but no usage mechanism or navigation November 23, 2026 (2026-11-23)

02 **ASM risk fingerprint for nginx2root** Not checked yet NEW

30

Censys added a new ASM (Attack Surface Management) risk fingerprint for the nginx2root vulnerability class.

— Mentioned only in summary, no detail on scope or usage November 23, 2026 (2026-11-23)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Knocknoc

Feedback ⓘ v26.11 NOTES ↗ RANK

Knocknoc is a just-in-time network access tool that grants users temporary firewall access only after authentication, reducing attack surface exposure.

> Knocknoc's v26.11 release adds resilient Agent sets with automatic failover, identity-based bastion egress control, self-service access revocation, and longer offline Enterprise licences, alongside SAML sign-in diagnostics and auto-recovery, SBOM visibility, and new API key authentication for machine-to-machine access.

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

01 **SBOM and multi-server log access on Support page** Not checked yet

NEW 70

Adds the ability to view or download a Software Bill of Materials (SBOM) for the Server and Agent , and shows recent logs from every server in multi-server deployments, both from the Support page.

Review and download the SBOM for the Server and Agent you are running to check component versions against your own vulnerability tooling.

> 📍 In the Knocknoc admin UI, go to Support, then locate the SBOM section to view or download the Software Bill of Materials for your current Server and Agent versions.

— Names exact artefacts and page, with example steps to access v26.11

02 **Agent sets with automatic failover** Not checked yet NEW

67

Multiple Agents can now be added to a single integration so access changes continue when one Agent's host is patched, rebooted, or offline. Admins control the failover order and each standby Agent gets a pre-save connectivity check.

— Explains mechanism and admin controls but no exact UI path v26.11

03 **SAML sign-in session recording for diagnostics** Not checked yet

NEW 64

Admins can temporarily record SAML sign-ins and inspect, per session, which details the identity provider sent and which groups or admin access they produced, making it easier to diagnose wrong-group membership.

Diagnose why a user landed in the wrong group after SAML sign-in by temporarily recording sign-in sessions and inspecting what the identity provider sent.

> 📍 In the Knocknoc admin UI, go to the SAML configuration section, enable temporary sign-in recording, have the user sign in, then open the recorded session to see which attributes and groups arrived from the identity provider.

— Usage example gives exact steps to diagnose sign-in issues v26.11

---

04 **API key and EDL secret addresses lose expiry support**

Not checked yet 63

BREAKING

API key and EDL secret addresses no longer support expiry dates; at upgrade, addresses that have already expired are removed and the rest become permanent. Admins should review any addresses used for temporary access before upgrading.

— States exact upgrade behaviour and a pre-upgrade action v26.11

---

05 **SAML encrypted assertions now require RSA-OAEP** Not checked yet

60

BREAKING

SAML with encrypted assertions now requires RSA-OAEP; older encryption settings are no longer accepted, so identity provider encryption settings should be checked before upgrading.

— Names exact cryptographic requirement and pre-upgrade check v26.11

THINNER COVERAGE BELOW

06 **Identity-based bastion egress control** Not checked yet NEW

58

Outbound connections from shared Linux, Windows, and OpenBSD bastion hosts can now be controlled per-user, so destinations open only for users who have been granted access and close again when access ends.

— Describes behaviour and scope but no config surface named v26.11

---

07 **Self-revoke access from any Knoc** Not checked yet NEW

49

Users can now revoke their own access from any Knoc — including Knocs that grant access automatically at sign-in — without waiting for expiry or signing out.

— Clear behaviour described, no exact UI location given v26.11

---

08 **90-day offline Enterprise licences** Not checked yet IMPROVED

45

Enterprise licences now support 90-day offline validity, with the extended period requiring no configuration changes after upgrading.

— States the new limit but little mechanism beyond that v26.11

09 **API key authentication for machine-to-machine access**

Not checked yet NEW 45

Authenticate to Knocknoc with an API key for machine-to-machine access, for use with the `knocknoc-client`.

— Names the client and auth method but no further detail product docs

10 **SAML sign-in auto-recovery from interruptions** Not checked yet

IMPROVED 37

SAML sign-in now recovers automatically from interrupted sessions — MFA delays or tab switches that previously showed error pages now restart the sign-in flow instead.

— Describes before/after behaviour but no config or flag v26.11

11 **Test a Knoc against its device from overview page** Not checked yet

NEW 35

Testing a Knoc against its device is now available directly on the Knoc's overview page.

— Names exact UI location but no further mechanism v26.11

12 **IP intelligence degraded status in Health panel** Not checked yet

NEW 30

The sidebar's Health panel now shows an IP intelligence degraded status when connectivity issues are detected.

— Brief description of a single status indicator product docs

└──▷ **BREAKING ON UPGRADE**

! API key and EDL secret addresses no longer support expiry dates; at upgrade, addresses that have already expired are removed and the rest become permanent — review any addresses used for temporary access before upgrading.

! SAML with encrypted assertions now requires RSA-OAEP; older encryption settings are no longer accepted — check your identity provider's encryption setting before upgrading.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

iDenfy

Feedback ⓘ updated · Nov 24, 2026 NOTES ↗ RANK

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy's latest update focuses on operational control over verification sessions and KYB workflows, adding session cancellation, partial form progress saving, a new expiry webhook, and expanded bank card and reporting checks.

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

01 **Bank card expiry and client ID linking** Not checked yet IMPROVED

45

Bank card verification now checks card expiry dates and links standalone bank card results to your own client ID.

– Describes two concrete checks with no mechanism detail snapshot-20261124

02 **Cancel unstarted verification sessions** Not checked yet NEW

40

Unstarted verification sessions can now be cancelled either through the API or directly from the dashboard.

– Names both surfaces but no endpoint or exact UI path snapshot-20261124

03 **KYB form partial progress saving** Not checked yet IMPROVED

40

Business verification (KYB) forms now save progress starting from the first field filled, allowing follow-up with clients who did not complete the form.

– Explains behaviour but no config or trigger details snapshot-20261124

04 **KYB link-expiry webhook** Not checked yet NEW

40

A new webhook reports when a KYB verification link expires.

– Names the webhook but no payload or setup details snapshot-20261124

05 **AI reviewer results in KYB PDF reports** Not checked yet NEW

25

KYB PDF reports can now include AI reviewer results.

– Bare statement with no detail on content or format snapshot-20261124

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Support for Personal Agent Protocol (Poppy) is now in Beta (2026-11-23) BLOG ↗

RANK

Secure users, AI agents, and more with Auth0, a production-ready authentication and authorization platform built to help you ship faster and drive sign-ups.

> Auth0 introduces Beta support for the Personal Agent Protocol (Poppy), letting users grant personal agents scoped, revocable access without sharing passwords.

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

01 **Personal Agent Protocol (Poppy) support** Not checked yet NEW

80

Auth0 now secures APIs and web apps for personal agents using the Personal Agent Protocol (Poppy), built on OAuth, JWTs, and DPoP, so users can grant scoped, revocable agent access while keeping passkeys and MFA, without sharing passwords. Configure personal agent support from the Auth0 CLI and verify setup with `auth0 poppy doctor`.

Verify your Poppy configuration is correct after setting up personal agent support in Auth0.

```
$ auth0 poppy doctor
```

— Names protocol, mechanism, and exact CLI command; Beta status and scope limits unclear

Support for Personal Agent Protocol (Po…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Bedrock Security

Feedback ⓘ announced · Nov 23, 2026

PRESS ↗ RANK

Bedrock Security is a more than DSPM: Bedrock's AI and data security platform unifies data, identity, and AI context so teams see real risk — not noise — without moving data.

> Bedrock Security launched Agent DLP, a runtime system that infers an AI agent's intent and evaluates its behavior against policy, backed by support for major AI assistants and gateways, an agent inventory with audit trail, and natural-language policies enforced against existing DSPM data classifications.

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

01 **Natural-language policy enforcement against DSPM classifications**

Not checked yet NEW 70

Security teams can write Agent DLP policies in natural language; Bedrock enforces them against Bedrock Data's existing DSPM classifications covering PII, PHI, payment card data, credentials and secrets, financial records, MNPI, and proprietary source code across on-premises, SaaS, and cloud.

— Names policy mechanism and classification categories, lacks exact syntax.

launch-20261123-8e775665

THINNER COVERAGE BELOW

02 **Runtime intent inference and behavior evaluation** Not checked yet NEW 55

Agent DLP infers the business purpose (intent) behind a user's task and evaluates the full session of an AI agent's actions against that intent and org policy at runtime, alerting, masking data, or blocking the agent when its behavior puts sensitive data at risk. It is available as a standalone product and as part of the Bedrock Data ArgusAI platform.

— Explains enforcement mechanism but no config/API surface given.

launch-20261123-8e775665

03 **Governed AI assistants and coding agents** Not checked yet NEW 50

Agent DLP now supports Claude, Claude Code, Claude Cowork, ChatGPT, Codex, Snowflake Cortex, and Snowflake Cortex Code as governed AI assistants and coding agents, applying intent-based policy enforcement to each.

— Names seven supported platforms but no setup steps. launch-20261123-8e775665

04 **AI gateway and agent runtime integrations** Not checked yet NEW

45

Agent DLP integrates with AI gateways and agent runtimes including Bifrost, LiteLLM, Amazon Bedrock AgentCore, and NVIDIA OpenShell.

— Names four integrations with no configuration detail. launch-20261123-8e775665

05 **Agent inventory and audit trail** Not checked yet NEW 45

Agent DLP inventories all agents in use along with the tools, MCP servers, and identities each relies on, and keeps a full audit trail of every action for regulatory review.

— Describes inventory/audit scope without a UI or API path. launch-20261123-8e775665

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Druva Data Resiliency Cloud Feedback ⓘ

changelog · Nov 23, 2026 NOTES ↗ RANK

Druva Data Resiliency Cloud is a SaaS platform that provides backup, disaster recovery, and data protection for endpoints, cloud workloads, and SaaS applications.

> Druva Data Resiliency Cloud introduces a unified dashboard for monitoring risk across backup workloads, infrastructure, and SaaS applications.

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

01 **Unified risk dashboard across workloads** Not checked yet NEW

35

A single dashboard now lets users view potential risks across backup workloads, infrastructure, SaaS applications, and endpoints in one place.

.

— Describes the dashboard's scope but no mechanism or navigation detail

2026-11-23-1161e433

WAS THIS USEFUL? 👍 👎

Druva inSync

Feedback ⓘ changelog · Nov 23, 2026 NOTES ↗

RANK

Druva inSync is a cloud-based data protection and governance platform that backs up endpoint and SaaS application data while ensuring compliance and security.

> Druva inSync introduces a unified dashboard consolidating risk visibility across backup workloads, infrastructure, and SaaS applications and endpoints.

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

01 **Unified risk visibility dashboard** Not checked yet NEW

35

A new dashboard surfaces potential risks across backup workloads, infrastructure, and SaaS applications and endpoints in a single view.

— Names scope but no UI path, metrics, or mechanism given. 2026-11-23-a63758a9

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

Riot

Feedback ⓘ updated · Nov 24, 2026 NOTES ↗

Riot protects 2M+ employees from phishing, deepfakes and social engineering with AI security awareness training, phishing simulations and email security.

Riot adds phishing-specific remedial training, a course preview tab with per-employee view as selection, and an Ask Albert Most asked questions admin view.

UPDATED INTEGRATION Admins can now see the security questions employees ask Albert most often, with anonymous real examples grouped by recurring question, in the Most asked questions view in Ask Albert in Inbox. Not checked yet

Remedial training after a phishing simulation now includes a page specific to how the employee was tricked — credential submission, MFA code entry, QR code scan, OAuth consent, or file open — when follow-up training is enabled in campaign settings. Not checked yet

NEW DETECTION Adds a Preview tab to the course panel so admins can see exactly what each employee receives on Teams, Slack, Google Chat, or the Portal, including delivery language, personalised images, and conversation path, using the view as selector. Not checked yet

NEW DETECTION Sonar rule: over-exposed LinkedIn profiles

UPDATED DETECTION Data breach monitoring moved into Sonar as a native rule

NEW COMPLIANCE FRAMEWORK OWASP Top 10 courses: Cryptographic Failures (#4) and Injection (#5)

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Mend.io Mend AI

Feedback ⓘ changelog · Nov 24, 2026

CHANGELOG ↗ RANK

Mend AI is a traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend AI added a scan history view for tracking AI risk trends across scans and introduced autocomplete with free-text fallback for provider fields in AI Analysis workflows.

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

01 **AI scan history and trend view** Not checked yet NEW 60

Mend AI adds an AI scan history view to track risk trends over time, compare findings across scans, and drill down into individual scan results, with accurate partial and failed status reporting, violation counts, and branch metadata alignment.

Review AI risk trends across multiple scans to spot regressions after a model or dependency change.

📍 In the Mend AI platform, go to AI scan history, select the target repository and date range, then drill down into an individual scan result to inspect violation counts and branch metadata.

— Describes mechanism and scope but no API or exact UI path given.

changelog-20261124-889150d5

02

Autocomplete for AI Analysis workflow provider fields

Not checked yet IMPROVED 60

AI Analysis workflows now support autocomplete with a free-text fallback for the Model Provider and Model Inference Provider author fields, allowing any detected vendor or a manually entered value to be used in workflow rules.

Create a workflow rule for a provider not in the predefined list by typing its name directly into the author field.

> 📍 In AI Analysis workflows, open the rule editor, click the Model Provider or Model Inference Provider author field, type the vendor name to trigger autocomplete or enter a custom value using the free-text fallback, then save the rule.

— Names exact fields and behavior, with UI steps but no API/flag.

changelog-20261124-889150d5

WAS THIS USEFUL? 👍 👎

◆ AI Red Teaming & Testing

NVIDIA SkillSpector

Feedback ⓘ v2.13.7 NOTES ↗ RANK

SkillSpector is a security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

> SkillSpector's latest release deepens scan configurability and detection coverage, adding YARA rule support to baseline generation, broader Python dependency parsing, and a new regional AI provider.

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

01 **YARA rules in baseline generation** Not checked yet IMPROVED

65

The `--yara-rules-dir` flag can now be used during baseline generation, allowing scans and baselines to apply the same additional YARA rules.

— Names exact flag and scope but no further mechanism detail v2.13.7

02 **MCP include_report response parameter** Not checked yet NEW

65

MCP callers can now pass `include_report=false` to omit the rendered report from the tool response.

— Exact parameter and effect given, minimal further context v2.13.7

03 **Dependency parsing for Poetry, PDM, Hatch, uv** Not checked yet

IMPROVED

60

SkillSpector now reads dependency declarations from Poetry, PDM, Hatch, and uv tables in `pyproject.toml`.

— Names all four tools and the config file, but no usage steps v2.13.7

THINNER COVERAGE BELOW

04 **MiniMax provider and updated NV Build default model**

Not checked yet NEW

55

SkillSpector adds MiniMax as a regional provider and updates the NV Build default model to GLM-5.3 Flash.

— Names provider and model version, no configuration steps given v2.13.7

05 **MITRE ATLAS tagging for tool-poisoning findings** Not checked yet

IMPROVED

45

MCP tool-poisoning findings are now tagged with MITRE ATLAS `AML.T0110`.

— Names the tag but no detail on how it's surfaced or used v2.13.7

WAS THIS USEFUL? 👍 👎

◆ Model & Agent Protection

Minimal Software Research minimal

Feedback ⓘ

release-bdc02696 NOTES ↗ RANK

Minimal is a secure, declarative runtimes for Claude Code, Gemini, Codex, Pi and more, so you can run multiple agents safely on your stack from a verified software supply chain.

> Minimal added secret management, deeper box audit and spec introspection, network/egress controls on task and box runs, and several diagnostics and auth visibility improvements, while removing two legacy commands in favor of new equivalents.

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

01 Box audit trail and spec authority listing

Not checked yet NEW

80

`min box audit` reads a box's audit trail, and the `--parent` flag merges in the parent box's audit records to trace full egress history. `min box spec` lists each store reference's authorities for a box.

Inspect a box's proxy-decision audit trail and merge its parent box's records to trace the full egress history.

```
$ min box audit mybox --parent
```

Export a box's current spec as JSON to pipe into other tooling or store as a snapshot.

```
$ min box show mybox -o json
```

— Names exact commands and flag with a runnable example. release-bdc02696

THINNER COVERAGE BELOW

02 Secret management via `min secret`

Not checked yet NEW

50

A new `min secret` command lets users store, list, and remove secrets.

— Names the command but not subcommand syntax or storage mechanism. release-bdc02696

03 Network and egress flags on task and box run

Not checked yet

NEW

50

`min task run` and `min box run` now take network and egress flags.

— Names commands but not the specific flag names or values. release-bdc02696

---

04 **Vulnerability reporting via `min diag --security`** Not checked yet

NEW

50

Users can send a suspected vulnerability report using `min diag --security`.

— Exact command named, but reporting process unexplained. release-bdc02696

---

05 **Box name completion for `min box` commands** Not checked yet

IMPROVED

45

Tab completion for box names is now available for every `min box` verb that accepts a box name.

— Clear scope but no concrete usage shown. release-bdc02696

---

06 **Edited audit chain detection in `min doctor`** Not checked yet

IMPROVED

35

`min doctor` now reports edited audit chains.

— Bare description of detection without mechanism. release-bdc02696

---

07 **Software key naming in `min auth status`** Not checked yet IMPROVED

35

`min auth status` now names each software key.

— Minor display change with no further detail. release-bdc02696

---

08 **Daily update check in `min`** Not checked yet NEW

35

The `min` CLI now checks once daily whether the channel has a newer release available.

— Simple behavior described without configuration options. release-bdc02696

└──▷ **BREAKING ON UPGRADE**

! The bare `min stop` and `min task run --path` commands are removed. Use `min host stop` and `-C <dir>` instead.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

PyPA python-pipx

Feedback ⓘ 1.19.0 NOTES ↗ RANK

The python-pipx CLI installs and runs Python applications in isolated environments while exposing their commands on the user's PATH.

> pipx 1.19.0 adds shell completion status reporting to `pipx list`, helping users spot missing or misplaced completion files for installed environments.

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

01 Shell completion status in `pipx list`

Not checked yet IMPROVED

70

`pipx list` now shows each environment's shell completions and flags missing or misplaced completion files, letting users audit completion setup across all installed environments.

Audit all installed pipx environments to confirm their shell completions are present and correctly placed.

```
$ pipx list
```

— Names exact command and behaviour change, clear to run 1.19.0

WAS THIS USEFUL? 👍 👎

goshs Feedback ⓘ v2.1.8 NOTES ↗ RANK

Goshs is a feature-rich single-binary file server for red teamers and developers. HTTP/S · WebDAV · FTP/SFTP · SMB · LDAP/S · NTLM hash capture · DNS/SMTP callbacks · TLS · Auth · Share links.

> goshs adds a flag and config key to run as a pure multi-protocol server with no HTTP listener, useful when operators want only FTP, SMB, LDAP or DNS exposure during an engagement.

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

01 **Run goshs without an HTTP listener** Not checked yet NEW

91

Passing `-nh` or `--no-http`, or setting `"http_server": false` in the config file, starts goshs with only its non-HTTP protocols (FTP, SMB, LDAP, DNS, ...) and no web server. The TUI, reverse-shell catcher, and `--user` privilege drop all continue to work, but `--tunnel` is refused in this mode and goshs refuses to start if no other protocol is enabled.

Run only FTP and LDAP listeners during an engagement where you want no HTTP exposure at all.

```
$ ./goshs -nh -ftp -ldap
```

— Names exact flag, config key, and documented restrictions/behaviour. v2.1.8

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Breachrx

Feedback ⓘ announced · Nov 24, 2026 PRESS ↗ RANK

BreachRx is an agentic AI CIRM platform that elevates cyber incident response to a governed enterprise-wide process.

> BreachRx launched Mobile Command, a mobile app for iOS and Android that brings incident response dashboards, reporting and secure messaging to executives outside the core platform.

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

01 **Mobile Command app for executives** Not checked yet NEW

50

BreachRx Mobile Command is now available on iOS and Android, giving C-suite executives live incident dashboards with open incidents, key metrics and completion progress, condensed non-technical summaries, and secure messaging outside the core platform. It adds real-time push notifications, chat and alerts, one-tap access to ongoing incident reports and centralized documents, and integrates with the full BreachRx platform so response teams can coordinate with executives in real time without context switching.

— Describes app features in prose but no API, config key or exact navigation.

launch-20261124-27354579

WAS THIS USEFUL? 👍 👎

HOL Guard

Feedback ⓘ v3.61.0 NOTES ↗

HOL Guard is an open-source runtime guard for AI agents that blocks risky tools, secret access, prompt injection, malicious packages, MCP servers, plugins and skills.

HOL Guard now runs the CIGAR engine natively via `guard-cigar-native` with no Python needed, adds opt-in git hook-bypass and whole-tree staging rules, and asks Codex for shorter replies when CIGAR is active.

└──▷ GET THIS VERSION ▾

pipx

```
$ pipx install hol-guard==3.61.0
# already installed? upgrade to this version:
$ pipx install --force hol-guard==3.61.0
```

git

```
$ git clone --branch v3.61.0 https://github.com/hashgraph-online/hol-guard.git
# already have the repo? check out this version:
$ git checkout v3.61.0
```

PyPI attestation · checked 24 Nov Wrong command?

Run the CIGAR engine natively inside Guard via the new `guard-cigar-native` crate, eliminating the Python dependency and serving the bridge (`guard-cigar-bridge.v1`), graph and broker workers (`cigar.context-worker.v1`), and per-agent broker client. Not checked yet

Add a reserved `builtin` installation ID that starts the CIGAR engine from Guard's own executable with no downloads or stored files. Not checked yet

NEW DETECTION Add opt-in `git_hook_bypass` rule that blocks skipping pre-commit and pre-push hooks via `--no-verify` / `-n` , `HUSKY=0` , `LEFTHOOK=0` , and related bypasses. Not checked yet

NEW DETECTION Add opt-in `git_stage_all` rule that requires staging files by name and catches `git add -A/-u/--all/.` , `:/` , :(top), `git commit -a/--all` , `git commit .` , `git rm -r --cached .` , and exclusion-only pathspecs. Not checked yet

Ask Codex for short replies when CIGAR is active. Not checked yet

NEW INTEGRATION AIfa Memory hosted MCP server

NEW INTEGRATION KinetAios MCP server

NEW INTEGRATION ForgeLoopBridge MCP server protection

NEW INTEGRATION Mercado Pago MCP protection

NEW INTEGRATION AgentBurp MCP protection

NEW INTEGRATION Server Guardian MCP protection

NEW INTEGRATION birkin-mnemosyne MCP server

NEW INTEGRATION Rozee job search MCP protection

NEW INTEGRATION Apple Docs MCP protection

NEW INTEGRATION Webcrawl MCP protection

NEW INTEGRATION agentic-linkedin MCP server

NEW INTEGRATION deed.page hosted MCP server

NEW INTEGRATION Zambo hosted MCP server policy

NEW INTEGRATION DocuSense MCP server

NEW INTEGRATION Lienzo UI MCP server

NEW INTEGRATION krusch-skylight MCP server and trust binding

NEW INTEGRATION hydrascope MCP server and trust binding

NEW INTEGRATION deadbolt MCP server and trust binding

NEW INTEGRATION krusch-frontend-mcp server and trust binding

NEW INTEGRATION krusch-fs MCP server and trust binding

NEW INTEGRATION krusch-grounded-harness MCP server and trust binding

NEW INTEGRATION krusch-stack MCP server and trust binding

NEW INTEGRATION LocalSites MCP server

NEW INTEGRATION New Jazz City chain reads MCP server

NEW DETECTION FlashPilot firmware write review

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 24, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue