THE CYBER TOOLCHAIN

NO. 145

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 20, 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 106 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 41 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. Contrast Security shipped a runtime block for Java4Shell (CVE-2026-114938) that stops the exploit at

vim ? help

Base64 decoding, buying teams cover until the JDK patch lands. GitLab now triages and remediates vulnerabilities across a whole group and auto-rotates leaked tokens. MIND can block outbound email at Exchange Online, closing the gap left by endpoint and browser DLP.

PROTECT

Block an unpatched Java exploit at runtime instead of waiting on the JDK

Contrast's Java agent now has a block rule for Java4Shell (CVE-2026-114938) that intervenes at the Base64 decoding stage, which is the only lever most teams have while a JDK fix is still upstream and emergency redeploys are off the table. Trend Vision One adds FIPS mode control on the Service Gateway appliance via CLISH and enforces VPC placement for File Security's AWS Lambda functions, so scanning workloads stop having a default route to the internet.

Contrast Security · Trend Vision One

PROTECT

Hand group-wide vulnerability triage and token rotation to an agent

GitLab's wave moves triage and remediation from per-project to group scope and rotates leaked tokens automatically — the window between a secret hitting a commit and someone noticing is where most credential incidents live. Endor Labs adds `endorctl` commands for scanning and remediating container base images, which is where a large share of inherited CVEs actually sit, plus SBOM/VEX evidence for FedRAMP.

GitLab · Endor Labs

PROTECT

Stop sensitive data on the way out of the mail server, not just the endpoint

Real-time outbound DLP blocking for Microsoft Exchange Online adds email as an enforcement point alongside endpoint and browser. Email has always been the leakiest channel and the hardest to cover from the endpoint alone, since webmail and mobile clients bypass agent-based controls entirely.

MIND

GOVERN

Query OpenShift compliance posture in plain English instead of reading Compliance Operator output

> Elastic polls the OpenShift Compliance Operator APIs into Elasticsearch, builds latest-state indices and dashboards, and puts an Agent Builder agent on top — compliance scan results stop being a point-in-time artifact someone exports for an auditor. Sublime scopes custom roles to specific message verdicts within each flag category, so a junior analyst can work malicious mail without seeing everything the detection stack flagged.
>
> Elastic Security Labs

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ updated · Nov 20, 2026 NOTES ↗ RANK

Hacktron is a review, Whitebox, the API, and MCP.

> Hacktron adds a signed Webhook notifier for automations, a configurable pull-request remediation workflow for findings, multi-application repository membership, and a cluster of refinements to Whitebox cost estimates, scan details, PDF reports and findings-list controls.

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

01 Webhook notifier for automations

Not checked yet NEW

85

Automations gain a Webhook notifier that posts signed JSON events to an HTTPS endpoint you configure, with deliveries recorded in run history; new-vulnerability automations can now notify without requiring an attached action.

Route new-vulnerability alerts from an automation directly to your incident-response pipeline without attaching a remediation action — useful for read-only notification workflows.

📍 In the automation editor, open the Settings tab, add a Webhook notifier, paste your HTTPS endpoint URL, and save. Under trigger, select New vulnerability — no action is required. Deliveries will appear in run history.

— Named mechanism (signed JSON, HTTPS endpoint) with exact setup steps snapshot-20261120

02 Remediate workflow with PR or patch default

Not checked yet NEW

80

The findings list gets a Fix column showing View fix when a fix is in review or Remediate on hover for open findings without one. Clicking Remediate lets an organisation admin choose Create a pull request or Prepare the fix only and save that choice as the organisation default for future runs.

Let org admins standardise whether engineers get a ready-to-merge pull request or a local patch when they click Remediate on a finding.

📍 As an organisation admin, open a finding and click Remediate, choose Create a pull request or Prepare the fix only, then save as the organisation default for future runs.

— Exact UI labels and admin steps to set the default snapshot-20261120

THINNER COVERAGE BELOW

03 **Delete finished Whitebox scans** Not checked yet NEW 55

Finished Whitebox scans and their reports can now be deleted via Delete scan in the scans list; the scan's findings are retained in Findings.

– Named UI action with clear retention behaviour snapshot-20261120

04 **Inactive findings stay visible with undo** Not checked yet IMPROVED 55

Findings moved to an inactive status now remain visible but dimmed in the findings list, with an Undo option shown in a Moved to inactive toast; the findings list no longer applies a default Status filter.

– Describes the toast and filter change clearly snapshot-20261120

05 **Whitebox cost estimates by code size** Not checked yet IMPROVED 50

Whitebox cost estimates now quote the whole pentest by code size, showing each repository's lines of code with an expandable folder breakdown of Lines and Credits.

– Names UI elements but no navigation path given snapshot-20261120

06 **Multi-application repository membership** Not checked yet NEW 45

Repositories can now belong to several applications at once; repository pickers, including Add context, display every application a repository belongs to, with a +N badge on hover.

– UI behaviour described without exact navigation snapshot-20261120

07 **Default organisation on sign-in** Not checked yet NEW 45

Personal settings now include Set as default to choose which organisation opens after sign-in.

– Named setting but minor scope snapshot-20261120

08 **Code scope panel in scan details** Not checked yet NEW 40

Scan details now include a Code scope panel showing each repository's include and exclude rules.

– Names the panel but no example of use snapshot-20261120

09 **Paths in scope table in PDF reports** Not checked yet IMPROVED 40

Full and Executive PDF reports now include a Paths in scope table and use the same finding IDs as the app.

— Describes the change but no interaction steps `snapshot-20261120`

WAS THIS USEFUL? 👍 👎

RedAmon

Feedback ⓘ v7.9.0 NOTES ↗ RANK

RedAmon is an open-source, self-hosted AI penetration testing framework: maps your attack surface into a graph, autonomously exploits it from a Kali sandbox with human approval gates, and opens PRs that fix what it finds. MCP both ways: plug in any MCP server as a tool, or drive RedAmon from Claude Code or your own agent.

> RedAmon's v7.9.0 release adds a project-aware AI Navigator chat assistant and overhauls Tradecraft into a 2.0 per-account knowledge library with three new agent tools, replacing 1.x in a breaking migration, while also shipping MCP tool-plugin presets, a stricter Neo4j APOC procedure set, and expanded accounts/roles documentation.

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

01 **Tradecraft 2.0 knowledge library and agent tools** Not checked yet 98

BREAKING

RedAmon introduces Tradecraft 2.0, a per-account knowledge library of trusted offensive-security sites the agent searches and reads mid-attack chain, managed under Global Settings → Tradecraft. Sites can be added via 105 presets or custom URL ingestion through Check URL, using Live (site-map and live page reads) or Stored (cleaned page copies in your database) ingestion modes, with site logins and a Map view to browse, search, pin, block, rename, and restore pages; Stored pages count against a quota of 1 GB per account and 20 GB per install, with 95 configurable ingestion and search settings. Three new agent tools query the library: `tradecraft_search` (keyword search with no model call), `tradecraft_lookup` (reads one section of a page by its page ID, e.g. `patt:412`), and `tradecraft_ask` (a librarian returning up to three checked pages and exact excerpts for a vague need). It also supports GitHub-token-authenticated CVE archive indexing (up to 20,000 largest CVE files, no text stored), docs-site mapping from search indexes or sitemaps, rule-based, LLM-guided, or Auto website crawls, and a focus mode (keep everything, Focused on chosen topics, or Tag only), with resources enabled per project via the Agent Skills & Tradecraft tab. This replaces Tradecraft 1.x for every account with no option to keep 1.x — resources convert automatically on first agent start after update; the old page-cache volume `redamon_tradecraft_cache` is deleted on update (run `touch .tradecraft-cache-removed` at the repository root beforehand to keep it); rolling back to 7.8.0 drops the Tradecraft 2.0 tables and data; projects that previously limited `tradecraft_lookup` to exploitation and post-exploitation phases now have it expanded to all three, logged as `system.tradecraft_phases_backfill`; and new projects start with the CVE (MSF) skill off by default.

Add HackTricks as a Tradecraft source so the agent can read the exact technique page it needs mid-attack chain.

📍 Go to Global Settings → Tradecraft, click `Check URL`, enter `https://book.hacktricks.xyz`, choose `Stored` ingestion mode, and save — the site is ingested in the background and the agent will surface relevant pages as `Library now` during attack steps.

Add a Stored Tradecraft resource and restrict it to specific topics so the agent draws only relevant pages during an engagement.

📍 Go to Global Settings → Tradecraft, add a new resource, set its type to Stored, set the ingestion focus to Focused, and select the topics to keep. Then open the project's Agent Skills & Tradecraft tab and enable the resource for that project.

— Names every config key, tool, quota and migration step with exact identifiers. v7.9.0

02 **Project-aware AI Navigator chat assistant** Not checked yet NEW

80

AI Navigator is a chat assistant available on every page that answers questions about how RedAmon works (linking to documentation) and about the selected project's data — including the graph, findings, scans, settings, triage, remediations, reports, and workspace files. Access is controlled via the new `navigator:use` permission, enabled or restricted instance-wide under IAM → Instance and granted per project through the share dialog; each user picks their own AI model at Global Settings → Models by feature → AI Navigator. Custom permission sets show `navigator:use` as `available, not granted`, so admins must manually add it to sets that should use Navigator, and rolling back to 7.8.0 drops every Navigator conversation along with every owner's collaborator choice.

Grant a collaborator access to AI Navigator on your project so they can query project findings and graph data via the chat.

📍 In the project share dialog, find the collaborator's entry and enable the Navigator toggle (requires the `navigator:use` permission to be present in their role). To allow or block Navigator for the entire instance, go to IAM → Instance and toggle the Navigator setting.

Pick the AI model that powers AI Navigator for your account so the chat can answer questions about your project data.

📍 Go to Global Settings → Models by feature → AI Navigator and select your preferred model. If no model is set, the first Navigator message will prompt you to choose one.

During a pentest, ask AI Navigator what findings are open in the current project without leaving your working context.

📍 In any RedAmon page, open the AI Navigator chat panel and type: 'What critical findings are in this project?' — it will answer using only the data you are permitted to read, citing the exact documentation section and project graph.

— Exact permission and settings path given, but model list and scope limits unstated.

v7.9.0

03 **MCP server tool plugins with presets** Not checked yet NEW

72

Any Model-Context-Protocol server can be plugged into the agent as a tool plugin, with 33 prefilled presets including Shodan, GitHub, Hugging Face, Censys, and mitmproxy. Each plugin offers a `Test/Discover` button to verify the connection and phase-aware Tool Matrix integration so it can be enabled for specific attack phases.

Connect a Shodan MCP server so the agent can query internet-exposed assets as a tool plugin during recon phases.

📍 Go to the agent's tool plugin settings, select the Shodan preset from the 33 prefilled Model-Context-Protocol presets, click `Test/Discover` to verify the connection, then enable it in the Tool Matrix for the Recon phase.

— Names presets and control, but no API schema or limits given. product docs

04 **Neo4j APOC procedure restriction** Not checked yet BREAKING

60

Neo4j now loads only the APOC procedures RedAmon calls, so scripts calling APOC procedures not used by RedAmon get a no procedure registered error — `apoc.load.*` and `apoc.cypher.*` are no longer available.

— Names exact removed procedure families but gives no migration path. v7.9.0

THINNER COVERAGE BELOW

05 **Accounts and roles documentation expanded** Not checked yet

IMPROVED

44

Documentation now covers all 99 permissions explained step by step, including signing in, inviting people, resetting credentials, disabling and deleting accounts, managing administrators, giving permission sets, and acting as a user.

— Documentation-only update with no new mechanism or command. product docs

└──▷ BREAKING ON UPGRADE

! Tradecraft 2.0 replaces Tradecraft 1.x for every account with no option to keep 1.x; resources are converted automatically on first agent start after update.

! Scripts calling APOC procedures not used by RedAmon will get `no procedure registered` — Neo4j now loads only the APOC procedures RedAmon calls, so `apoc.load.*` and `apoc.cypher.*` are no longer available.

! Rolling back to 7.8.0 drops every AI Navigator conversation, every owner's collaborator choice, and the Tradecraft 2.0 tables and columns with their data.

! New projects now start with the CVE (MSF) skill off by default.

! The old Tradecraft page-cache volume (`redamon_tradecraft_cache`) is deleted on update; run `touch .tradecraft-cache-removed` at the repository root before updating to keep it.

! Custom permission sets show `navigator:use` as `available, not granted` — admins must manually add it to sets that should use AI Navigator.

! Where `tradecraft_lookup` was allowed in exactly exploitation and post-exploitation phases, it is automatically expanded to all three phases; the audit log's `system.tradecraft_phases_backfill` row lists affected projects.

WAS THIS USEFUL? 👍 👎

◆ **Post-Exploitation**

SpecterOps BloodHound

Feedback ⓘ v9.10.1 NOTES ↗

RANK

Bloodhound is a six Degrees of Domain Admin

> BloodHound added configurable limits for PostgreSQL graph traversals to control resource usage during analysis.

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

01 **Configurable traversal limits for PostgreSQL** Not checked yet

NEW 50

BloodHound now supports configurable limits for PostgreSQL graph traversals to control traversal state and byte usage, available via updated DAWGS v1.0.2.

— Names DAWGS version but no config key or flag given v9.10.1

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

OpenQodex

Feedback ⓘ openqodex@0.22.0 NOTES ↗ RANK

OpenQodex is an open source AI code review for coding agents. A separate reviewer checks every change your agent writes before you push, in Claude Code, Codex, Cursor and Cline.

> OpenQodex's 0.22.0 release expands its code graph with gin service support and deeper TypeScript resolution, while improving coverage-gap reporting and fixing import/reference accounting bugs across languages.

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

01 Gin service support in `graph routes`

Not checked yet NEW

90

`openqodex graph routes` now reads gin services, including engines, groups with their prefixes, `Use` middleware, and handlers; inline route handlers (e.g. app.get('/x', (c) => ...)) are treated as first-class definitions, and the new `--exclude-tests` flag omits routes whose test files build the application.

Audit only production routes in a gin-based Go service, skipping routes defined inside test files.

```
$ openqodex graph routes --exclude-tests
```

— Names exact command and flag with runnable example openqodex@0.22.0

02 Expanded TypeScript type reference resolution

Not checked yet

IMPROVED

83

`openqodex graph references` on a TypeScript type now finds uses in interface properties, type aliases, type arguments (e.g. new Map<string, Repo>()), type parameter constraints and defaults, and classes that implement or extend the type.

Find all call sites and implementors of a TypeScript interface, including structurally-matched classes and type-argument uses that were previously missed.

```
$ openqodex graph references ISessionStore
```

— Named command, scope, and runnable example provided openqodex@0.22.0

03 Coverage gap reporting for unread languages

Not checked yet NEW

78

Code graph answers now surface coverage gaps for unread languages — Rust, C, C++, C#, Kotlin, Swift, PHP, and Scala — with reasons like `2 Rust files not read`; each answer gains an `unknown.groups` field breaking gaps down by cause with a count and up to two examples, plus a `listIncomplete` flag, and the CLI prints `certain (list incomplete)` in the counts line.

— Names field and flag but no runnable example `openqodex@0.22.0`

04 **Per-import-line counting in `graph packages`** Not checked yet

IMPROVED

65

`openqodex graph packages` now counts every import line separately, so a file importing the same package via `import type` and `import` no longer loses the duplicate site; `importers` still lists each file once.

— Clear before/after but no usage example given `openqodex@0.22.0`

THINNER COVERAGE BELOW

05 **External classification for chained TypeScript calls**

Not checked yet IMPROVED

55

TypeScript method calls on values returned by a package are now classified as external at any chain depth, instead of raising an unknown gap that set the floor for every same-named repository method.

— Explains mechanism and prior bug, lacks concrete example `openqodex@0.22.0`

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Improvements on the pull request overview in public preview (2026-11-20) BLOG ↗

RANK

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

> GitHub expanded its public preview of the pull request Overview page, adding clearer merge-requirement and check signaling, force-push commit comparison links with timestamps, new sidebar actions for locking conversations and editing development links, and smaller UI and Copilot refinements.

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

01 **Clearer merge requirements and check previews** Not checked yet

IMPROVED 55

In the Overview page, failed optional checks no longer make mergeable pull requests appear blocked, previously missing merge requirements now show what needs attention, and checks and approvals now expand when clicked instead of pinning a preview.

— Describes behavior change but no config or flag to act on

Improvements on the pull request overvi…

02 **Force-push comparison links and commit timestamps** Not checked yet

IMPROVED 50

In the Overview page, force-push activity now links to the old and new commits so you can compare changes, and commit activity now includes timestamps so you can see when changes were made.

Opt in to the Overview page public preview to access the new sidebar actions and commit comparison links on your next pull request.

> 📍 In the pull request header, select the ✨ Try the Overview page button. To return to the classic view, use the same menu.

— Clear mechanism but only UI navigation, no command

Improvements on the pull request overvi…

03 **Sidebar actions for locking and development links** Not checked yet

NEW 35

The Overview page gained sidebar actions to lock or unlock PR conversations directly in the view, and to edit Development sidebar links without leaving the page.

— Named UI actions but no deeper mechanism given Improvements on the pull request overvi…

04 **Responsive branch display on Overview page** Not checked yet

IMPROVED 20

Source and target branches now stay visible on smaller screens in the Overview page.

— Minor UI tweak with no further detail Improvements on the pull request overvi…

05 **Copilot mention clarity before submission** Not checked yet

IMPROVED 20

Comments mentioning Copilot now better explain the action and selected model before submission in the Overview page.

— Brief description, no specifics on models or flow

Improvements on the pull request overvi…

NEW INTEGRATION Enterprise Live Migrations from Azure DevOps Services to GitHub Enterprise Cloud

NEW SUPPORTED PLATFORM RHEL 9 and RHEL 10 runner images generally available for GitHub-hosted larger runners

> Enterprise Live Migrations from Azure DevOps generally available (2026-11-19)
>
> BLOG ↗ STABLE
>
> GitHub adds live repository migration from Azure DevOps, syncing history, branch policies as rulesets, Boards connections, and Pipelines with continuous sync until cutover.
>
> Migrate repositories from Azure DevOps Services to GitHub Enterprise Cloud with data residency while developers keep contributing, then cut over at a time you choose. Not checked yet
>
> Manage migrations through the Azure DevOps portal for a guided experience or script them with the Azure DevOps CLI. Not checked yet
>
> Pre-migration validation and lifecycle tracking identify issues and monitor progress before cutover. Not checked yet
>
> Sync Git history, branches, tags, and pull requests — including titles, descriptions, comments, and user history — from Azure DevOps to GitHub. Not checked yet
>
> Automatically migrate Azure DevOps branch policies as GitHub branch rulesets. Not checked yet
>
> +1 more

> Red Hat Enterprise Linux runner images now generally available (2026-11-19)
>
> BLOG ↗ STABLE
>
> GitHub Actions adds RHEL 9 and RHEL 10 runner images as generally available for production workflows on Team and Enterprise Cloud plans.

WAS THIS USEFUL?  

cve-lite-cli

Feedback ⓘ v1.42.0 NOTES ↗

Cve-lite-cli is a fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

cve-lite-cli adds multi-format SBOM output, SBOM advisory database metadata, `--allow-empty` for empty-project scans, and output sanitisation against workflow command injection.

└──▷ **GET THIS VERSION** ▾

npm

```
$ npm install -g cve-lite-cli@1.42.0
```

git

```
$ git clone --branch v1.42.0 https://github.com/OWASP/cve-lite-cli.git
# already have the repo? check out this version:
$ git checkout v1.42.0
```

named in the project's README · checked 20 Nov Wrong command?

Request multiple SBOM formats in one run by passing `--sbom cyclonedx,spdx` or repeating the flag (`--cdx --sbom spdx`); each format is written to a separate file sharing one timestamp. Not checked yet

CycloneDX and SPDX SBOMs now record advisory database metadata: whether the scan ran offline, when the local database was last synced, and where it came from. Not checked yet

Allow empty-project scans without an exit code change by passing `--allow-empty` or setting the GitHub Action's `allow-empty` input. Not checked yet

NEW SUPPORTED PLATFORM Python requirements.txt scanning via OSV PyPI data

NEW INTEGRATION GitLab CI/CD template

└──▷ **BREAKING ON UPGRADE**

! Node 22.13.0 or later (22.x line) or 23.4.0 or later is now required; earlier Node versions are no longer supported.

! A scan that finds no lockfile and no exact-pinned `package.json` now exits 3 instead of 0, with or without `--fail-on` ; pass `--allow-empty` or set the Action's `allow-empty` input to restore the old behaviour.

- ! `cve-lite --report <dir>` with no project path is now refused; use `--report=<dir>` or put the project path before `--report`.
- ! Some CycloneDX severity ratings come out lower because a severity no longer carries over from one advisory to another for the same package.
- ! The first scan after upgrading runs with a cold OSV cache because the cache format changed.
- ! A known malicious package (MAL- advisory) now fails `--fail-on` at any level, including the default, where it previously passed every threshold except `low`.

WAS THIS USEFUL? 👍 👎

Contrast Security

Feedback ⓘ

Java4Shell (CVE-2026-114938) Analysis and Remediation (2026-11-19) BLOG ↗

RANK

Contrast Security is a secure AI-generated code with Contrast. Our runtime application security platform observes real behavior to block attacks in real time.

> Contrast Security shipped a runtime block rule for the Java4Shell exploit (CVE-2026-114938), letting the Java agent stop attacks at the Base64 decoding stage ahead of a JDK patch.

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

01 **Java4Shell runtime block rule for Java agent** Not checked yet

NEW 86

Contrast Java agent 6.38.0 or later adds a Java4Shell rule that can block or monitor CVE-2026-114938 exploit attempts at the Base64 decoding stage inside the JVM, stopping attacks before patching is complete. As an interim workaround on hosts where neither patching nor agent deployment is possible, the JVM flag `-XX:-UseBASE64Intrinsics` can be set (e.g. `java -XX:-UseBASE64Intrinsics -jar your-service.jar`), at a possible performance cost.

Put internet-facing Java services into block mode immediately while waiting for JDK vendors to ship fixed builds.

```
$ java -XX:-UseBASE64Intrinsics -jar your-service.jar
```

— Names agent version, rule, JVM flag and runnable command

Java4Shell (CVE-2026-114938) Analysis a…

WAS THIS USEFUL? 👍 👎

qodana-cli

Feedback ⓘ v2026.3.1 NOTES ↗ RANK

Qodana-cli jetbrains Qodana's official command line tool

> Qodana-cli's v2026.3.1 release adds an air-gapped scanning mode for internet-isolated environments and expands the Poly image's language coverage with Lua support and auto-detection.

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

01 **Air-gapped scanning mode** Not checked yet NEW 85

Enable air-gapped mode for environments with no internet access using the `--air-gapped` flag on `qodana scan`, or persist it project-wide via the `airGapped` field in `qodana.yaml` so every CI run scans offline automatically.

Scan a project in an air-gapped environment where the analysis container has no outbound internet access.

```
$ qodana scan --air-gapped
```

Persist air-gapped mode in your project config so every scan on any CI runs offline automatically.

```yaml
airGapped: true
```

— Names exact flag and config field with runnable examples v2026.3.1

THINNER COVERAGE BELOW

02 **Lua and Dart support in Poly image** Not checked yet NEW 45

Adds the Lua plugin to the Qodana Poly image, extending its language coverage, and adds auto-detection of Lua and Dart projects so the CLI now suggests the Qodana Poly linter when these are the dominant languages.

— Describes what changed but no command or config shown v2026.3.1

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 chunked sandbox tarball uploads and now retries Codex turns that end without a completed result

Uploads sandbox tarballs larger than 32 MiB in parts via `uploadTarballToSandbox`, writing each chunk as `<remoteTarPath>.part-NNNNN` and joining them with `cat` on the sandbox, so scans of large repositories no longer fail with `Request Entity Too Large` errors. Not checked yet

Codex batches whose turn ends without `turn.completed` — for example when OpenAI's safety filter blocks a turn mid-way — now retry and then fail instead of recording incomplete agent messages as clean findings. Not checked yet

WAS THIS USEFUL? 👍 👎

GitLab Feedback ⓘ

Cut your vulnerability backlog with agentic remediation (2026-11-19) BLOG ↗

RANK

Teams and agents ship trusted software faster with GitLab, the intelligent orchestration platform with context and governance for the entire software lifecycle.

> GitLab's November 19 wave is almost entirely about agentic engineering: group-wide vulnerability triage and remediation, a Slack-native agent, scheduled flow automations, a new Business Logic Security Analyzer, automatic leaked-token rotation, and a set of coding-agent skills for MR follow-up, code review, and issue-to-MR delivery.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **Leaked PAT auto-rotation via Secrets Manager** Not checked yet

NEW 95

Enable the rotation workflow on GitLab Self-Managed with the `secrets_manager_pat_rotation` feature flag (on by default on `GitLab.com`, not yet on GitLab Dedicated). The `leaked-token-rotation-automation` tool reads GitLab PATs — including fine-grained PATs — from CI/CD project variables, which secret detection does not scan, rotates them via the token's own `api` or `self_rotate` scope, stores the new values in Secrets Manager, and opens an MR for any jobs that reference the variable; cloud provider keys and other secret types are not yet covered.

Execute batch rotation across a group, capping at 25 rotations in one run to stay within credits budget.

```
$ leaked-token-rotation-automation --apply --limit 25
```

After rotation, the generated MR updates your pipeline so CI/CD jobs read the new secret from Secrets Manager at runtime instead of a hardcoded value.

```yaml
deploy:
  secrets:
    GITLAB_TOKEN:
      gitlab_secrets_manager:
        name: ROTATED_GITLAB_TOKEN_123
        file: false
  script:
    - ./deploy.sh
```

— Names flag, tool, scopes, and a runnable command

Rotate a leaked GitLab token and let Gi…

---

02 **Duo setup CLI skills and security triage-and-fix prompt**

Not checked yet NEW 90

The `gitlab-duo-setup` CLI now installs the `gitlab-mr-followup` skill, which checks open review threads, failed jobs, and conflicts on a merge request and works through them from your agent session, with both skills living in `~/.claude/skills/` or `.claude/skills/` inside the repo for team sharing and MCP connectivity verifiable via `/mcp` and the `list_duo_agents_and_flows` tool call; on GitLab Ultimate, `gitlab-duo-setup` also offers an optional security findings prompt that, when answered `yes`, enables the security triage-and-fix workflow and adds security flows to the group.

Bootstrap a top-level group for agentic engineering — installs the setup and MR-followup skills, then runs the guided setup from inside Claude Code.

```
$ git clone https://gitlab.com/gitlab-da/projects/agentic-eng-config.git
cp -R agentic-eng-config/skills/gitlab-duo-setup ~/.claude/skills/
cp -R agentic-eng-config/skills/gitlab-mr-followup ~/.claude/skills/
# then, inside Claude Code:
/gitlab-duo-setup
```

Bootstrap the full Duo agent workflow for your team so every MR gets automated review and pipeline repair by default.

> 📍 Install the skills from `https://gitlab.com/gitlab-da/projects/agentic-eng-config` and run gitlab-duo-setup for my GitLab group. Show me the plan before you change anything.

— Exact paths, skill names, and a verification command given Exact prompt and setup step named, little else

Agentic engineering with GitLab and you… | The best way to use your coding agents …

03 **Agentic @GitLab in Slack** Not checked yet NEW 85

Mention `@GitLab` in a Slack channel or DM to have an agent read the conversation, create issues and merge requests, and start agent flows, with results posted back to the thread; ask it to summarise a channel with `@GitLab, summarize #channel for me`. Each request runs as the requester's linked GitLab account, scoped to their existing permissions, creates an auditable agent session, and draws from GitLab Credits. It does not operate in external Slack Connect channels.

Turn a live Slack channel discussion into a tracked GitLab issue without leaving the thread.

> 📍 Steps: 1. Open the Slack channel where the discussion is happening. 2. Type '@GitLab create an issue from this thread' as a message. 3. The agent reads the conversation, creates the issue in GitLab under your account permissions, and posts a link back to the thread.

Catch up on a busy channel before asking the agent to act, so the resulting work item reflects the full discussion.

> 📍 Steps: 1. In any public or private Slack channel, type `@GitLab, summarize #channel for me`. 2. Review the summary the agent posts back. 3. Follow up with '@GitLab create a merge request for the agreed approach' to convert the decision into tracked work.

— Concrete commands, permission model, and scope limits given

Duo Agent Platform now does up to 2x mo… | Turn Slack conversations into trusted s…

04 **Scheduled automations for Duo Agent Platform flows**

Not checked yet NEW 85

Schedule Duo Agent Platform flows to run recurring tasks — from every 15 minutes to every month, at a chosen time and timezone — with no event trigger required (beta in GitLab 19.6). The redesigned Automations page (previously called `Triggers`) shows run volume, success rate, and session details for both scheduled and event-based automations. Each scheduled automation runs as a dedicated service account with scoped permissions, so agent actions appear separately from human actions in audit events, and every change to an automation is itself recorded in audit events.

Turn on experimental and beta features and set up a first scheduled automation — for example, an hourly CI health investigation — to verify the beta is active on your instance.

📍 In GitLab 19.6, go to your project › Automations (formerly Triggers) › New automation. Describe the task in plain language (e.g. 'Check open CI incidents and comment with the likely root cause'), set the cadence to `Every hour`, choose your timezone, and save. Confirm the automation appears with run count and success rate on the Automations page.

— Exact UI path, cadence range, and audit behaviour spelled out

Keep recurring work moving with schedul… | Duo Agent Platform now does up to 2x mo…

---

05 **Group-wide ATR profiles for vulnerability triage and remediation**

Not checked yet | NEW | 80

Apply automated triage and remediation (ATR) profiles across many projects at once from Security Inventory or the security configuration page, selecting by group or attribute, without scripting against the API. Profiles offer three presets — Conservative, Standard, or Proactive — covering false-positive detection, risk contextualization, and vulnerability resolution.

Roll out agentic triage and remediation to every project in a group without writing API scripts — pick a preset that matches your risk appetite.

📍 In GitLab Ultimate, go to Security › Security Inventory (or the security configuration page), select projects by group or attribute, choose an ATR profile preset (Conservative, Standard, or Proactive), and click Apply.

— Names presets and UI path but no API detail | Cut your vulnerability backlog with age…

---

06 **Vulnerability Context Flow exploitability risk table**

Not checked yet | NEW | 75

Enable Vulnerability Context Flow to add a risk table to SAST findings showing whether exploiting a flaw requires authentication or elevated authorization, whether the service handles sensitive data, and whether it is exposed to untrusted input, so teams can rank the backlog by exploitability as well as severity.

Add exploitability context to SAST findings so your team can prioritize public-facing services over isolated internal ones without manual triage.

📍 In GitLab Ultimate 19.6, follow the Vulnerability Context Flow setup to enable the feature; once on, open any SAST finding to see the new risk table showing authentication requirements, sensitive data exposure, and untrusted-input exposure.

— Clear mechanism and fields, setup path given | Cut your vulnerability backlog with age…

---

07 **Business Logic Security Analyzer (BLSA)** Not checked yet NEW

65

Self-service beta scans repositories and merge requests for broken ownership logic, missing authorization checks, multi-step abuse paths, and other authorization/access-control vulnerabilities that pattern-based scanners miss. Available via self-service on `GitLab.com` (GitLab Ultimate); Self-Managed and Dedicated customers must contact GitLab for access, and it can be powered by Anthropic's Claude Mythos 5.2 on request.

Catch broken ownership logic and missing authorization checks that static scanners miss by enabling Business Logic Security Analyzer on `GitLab.com` without waiting for an account team.

> 📍 In GitLab Ultimate on `GitLab.com`, follow the business logic scanning setup to self-enable BLSA; for Self-Managed or Dedicated, contact GitLab for access.

— Named scope and access paths, but still beta/thin mechanism

Cut your vulnerability backlog with age… | Duo Agent Platform now does up to 2x mo…

---

08 **GitLab Credits tracking, caps, and efficiency improvements**

Not checked yet IMPROVED

60

Prompt caching on the Duo Agent Platform lets eligible tasks use fewer model requests, stretching GitLab Credits up to 2x further; individual GitLab Credits consumption can now be tracked on `GitLab.com` filtered by product and date range without needing an administrator report; the Duo Agent Platform Impact dashboard (introduced in GitLab 19.5) shows adoption, work activity, and credit consumption broken down by model, flow type, and group or project; and a default per-user credit cap with individual overrides can be set on the Credit caps page (introduced in GitLab 19.4 for `GitLab.com` and Self-Managed) without using an API.

.

— States benefit but no config or toggle named Named pages and filters but no exact navigation steps

Duo Agent Platform now does up to 2x mo…

---

09 **/goal objective planning in Duo CLI and Chat** Not checked yet

NEW

60

Describe an objective with `/goal` in the Duo CLI or Duo Agent Platform Chat to have GitLab Duo plan, execute, and judge completion of the work, with the ability to pause, change, or redirect the agent at any point.

Hand off a multi-step objective to Duo so it plans, executes, and self-judges completion while you stay in control.

```
$ /goal Triage all open issues labelled 'needs-triage', assign severity, and add a comment summarising the action taken on each.
```

— Runnable command shown, but little on how judging works

Duo Agent Platform now does up to 2x mo…

10 **Developer flow: issue to merge request** Not checked yet NEW

60

Hand an assigned issue to the Developer flow and Duo runs in GitLab, reads the issue and repository, writes the change, and opens an MR in the user's name; governance and approval rules can still enforce a human review.

Hand a well-scoped issue to Duo's Developer flow so it writes the code and opens an MR without you copying context into a prompt.

> 📍 Find my assigned issues in <project>. Hand the best-scoped one to the Duo Developer flow.

Review Duo's pipeline fix and code review feedback on a specific MR before deciding what to merge.

> 📍 Did Duo propose a fix for the failed pipeline on MR !42? Show me the diff and why it should fix the failure.

— Clear workflow, but mechanism of code generation unexplained

The best way to use your coding agents …

THINNER COVERAGE BELOW

11 **Automated $0.25 second-look code review** Not checked yet NEW

50

Duo now gives every eligible merge request an automated second-look code review for a flat $0.25 per review when the MR is marked ready, posting feedback on the diff.

— Concrete price and trigger but no setup path The best way to use your coding agents …

12 **MCP tool calls recorded in Audit events** Not checked yet NEW

40

Tool calls made through GitLab's MCP server can now be viewed in Audit events when storage is enabled.

— Names the surface but gives no setup detail The best way to use your coding agents …

13 **Model choice per agentic flow** Not checked yet IMPROVED

40

Choose the model for each agentic flow from Anthropic, OpenAI, Google, or open-weight providers.

— Names providers but no selection mechanism shown

The best way to use your coding agents …

14 **Composite identity for agent actions** Not checked yet NEW

35

Composite identity now ties every agent action to both the agent and the person who started it, limiting access to what both are permitted to do.

— Describes behaviour only, no config or UI path The best way to use your coding agents …

WAS THIS USEFUL?  

◆ Software Supply Chain Security

Endor Labs

Feedback ⓘ

FedRAMP's new vulnerability rules: what changes for cloud providers | Endor Labs (2026-1

BLOG ↗ RANK

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC.

> Endor Labs shipped new `endorctl` commands for scanning and remediating container base images alongside SBOM/VEX evidence generation for FedRAMP compliance, and launched a deep integration with OpenAI's Daybreak Defense Factory that feeds reachability-based remediation guidance into Codex Security.

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

01 **Base image update scanning and remediation commands**

Not checked yet NEW 90

Scan a container and check for available base image updates with `endorctl container scan --base-image-check-updates` (plus `--namespace`, `--image`, and `--project-name`), then list candidate base image updates with `endorctl container remediation list-base-image-updates` (same arguments), which shows findings fixed and introduced by severity and reachability for each option.

Check a container image for available base image updates before shipping, to see which upgrade resolves the most reachable findings without introducing new ones.

```
$ endorctl container scan \
  --namespace=<your-namespace> \
  --image=<image_name:tag> \
  --project-name=<endor_project_name> \
  --base-image-check-updates
```

Review the ranked list of candidate base image updates — with per-severity and per-reachability counts of findings fixed and introduced — to pick the safest upgrade path.

```
$ endorctl container remediation list-base-image-updates \
  --namespace=<your-namespace> \
  --image=<image_name:tag> \
  --project-name=<endor_project_name>
```

— Exact commands and flags given, runnable immediately.

FedRAMP's new vulnerability rules: what…

02 **OpenAI Daybreak Defense Factory integration with Codex Security**

Not checked yet NEW 80

An AI SAST agent gives Codex evidence and remediation guidance for vulnerabilities in first-party source code; an SCA Remediation agent gives Codex reachability evidence and remediation guidance for open source vulnerabilities, cutting findings by up to 99% via full-stack reachability across direct and transitive dependencies; container base image vulnerabilities also get Codex reachability evidence and remediation guidance. A code context graph tells agents how the application uses each dependency, reducing agent token usage by 92% and completing security tasks 19x faster in Endor Labs benchmarks, and real-time vulnerability intelligence screens more than 50,000 open source packages per month for malicious behavior (including risks without a CVE), surfacing this data to Codex. Setup follows the OpenAI Daybreak Defense Factory import guide: open the console, select Endor Labs in the integrations section, authenticate, and findings with reachability, risk, and remediation guidance appear in Codex for automated fix implementation.

Follow the Daybreak Defense Factory import guide to connect Endor Labs findings to Codex Security for end-to-end automated remediation.

> 📍 Steps: 1. Open the OpenAI Daybreak Defense Factory console. 2. Navigate to the integrations section and select Endor Labs. 3. Follow the import guide to authenticate and configure the connection. 4. Endor Labs findings — with reachability, risk, and remediation guidance — will appear in Codex for automated fix implementation.

— Rich mechanism and numbers, but setup is a UI import guide, not a command.

Endor Labs + OpenAI: Strengthening the …

THINNER COVERAGE BELOW

03 **SBOM/VEX evidence and finding deduplication for compliance**

Not checked yet NEW 40

Endor Labs now deduplicates findings across SCA and container scans and generates SBOM and VEX evidence attached to every vulnerability verdict, producing machine-readable reporting intended for FedRAMP VDR/VER rules.

— Describes behavior but gives no command, UI, or config to act on.

FedRAMP's new vulnerability rules: what…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

openbao

Feedback ⓘ v2.6.6 NOTES ↗ RANK

OpenBao is a software solution to manage, store, and distribute sensitive data including secrets, certificates, and keys.

> OpenBao's latest releases add per-mount control over entity/group alias case-sensitivity and extend the Transit key-export API to support external-key public keys.

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

01 **Per-mount alias case-sensitivity control** Not checked yet

BREAKING

75

Auth mounts can now control alias case-sensitivity using the `case_sensitive_alias` config key. Entity and group aliases on newly created auth mounts are now case-sensitive by default, whereas previously they were case-insensitive; existing mounts are unaffected and can use `case_sensitive_alias` to override the new default.

– Names exact config key and before/after behaviour change v2.7.3 v2.6.6

THINNER COVERAGE BELOW

02 **Public key export for Transit external-key type** Not checked yet

IMPROVED

40

The existing key-export API can now export public keys for Transit keys of type `external-key`.

– Names API and key type but no endpoint path or usage detail v2.7.3

└──▷ **BREAKING ON UPGRADE**

! Entity and group aliases on newly created auth mounts are now case-sensitive by default, which may break setups that relied on case-insensitive alias matching.

! Entity and group aliases of newly created auth mounts are now case-sensitive by default (previously case-insensitive). Existing mounts are unaffected; use `case_sensitive_alias` to override.

WAS THIS USEFUL? 👍 👎

Truffle Security trufflehog

Feedback ⓘ v3.100.0

NOTES ↗ CODE ↗ RANK

TruffleHog finds, verifies, and analyzes leaked credentials across Git, chat, wikis, cloud storage, and more

> TruffleHog adds a Baseten API key detector with permission analysis, surfaces upsert errors from AWS Secrets Manager sources, and adds new metrics for detectors that exceed their timeout.

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

01 **Baseten API key detector with permission analysis** Not checked yet

NEW 81

Adds a Baseten API key detector, enabled via `feature.BasetenDetectorEnabled`, with permission analysis available through `baseten.AnalyzeAndPrintPermissions` using the `baseten` case in the analyzer CLI.

– Names exact flag, function, and CLI case v3.100.0

02 **Detector timeout overrun metrics** Not checked yet NEW

70

Adds metrics for detectors that run past their timeout: a `trufflehog_scanner_detector_timeout_ignored_total` counter, a `trufflehog_scanner_detector_workers_over_timeout` gauge, and an overrun-duration histogram, all labelled by `detector_name`.

– Names exact metric names and labels, no usage example v3.100.0

THINNER COVERAGE BELOW

03 **Upsert error field in AWS Secrets Manager source** Not checked yet

IMPROVED 50

Adds `upsert_error` to `AWSSecretsManager` source metadata fields, surfacing upsert failures in scan output.

– Names exact field but minimal mechanism detail v3.100.0

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

WPScan

Feedback ⓘ v4.2.1 NOTES ↗ CODE ↗ RANK

WPScan WordPress security scanner. Written for security professionals and blog maintainers to test the security of their WordPress websites. Contact us via contact@wpscan.com

> WPScan v4.2.1 adds visibility into responses truncated by `--max-response-size`, flagging incomplete scan data instead of silently missing it.

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

01 **Truncated response reporting** Not checked yet NEW 83

Responses cut short by `--max-response-size` are now listed in the scan summary: a count warning appears in the CLI output (with full URLs shown when `--verbose` is passed), and a `truncated_responses` array is added to JSON output, so findings based on incomplete bodies are flagged rather than silently missed.

— Names exact flags and JSON field; mechanism and purpose both explained. v4.2.1

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Wallarm

Feedback ⓘ Your security team just got bigger (2026-11-19)

BLOG ↗ RANK

Discover the latest features, improvements, and updates in Wallarm API Security

> Wallarm rolled out seven AI Agentic Operators that automate alert triage, rule hygiene, and policy tuning, manageable via its MCP server and with an optional autonomous execution mode.

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

01 **Seven AI Agentic Operators for triage and tuning** Not checked yet

NEW 65

Wallarm introduces seven Agentic Operators: Radar, which detects attack spikes, investigates them, and suggests a response; Sift, which clears false positives on an hourly cycle so alerts remain actionable; Trim, which tidies the active ruleset to prevent rule sprawl; Tune, which adjusts protection policies to match observed application behaviour; Lens, which reviews the API inventory and flags risky endpoints; Bard, which converts the current security posture into a human-readable briefing; and Wally, which answers questions and handles configuration requests on demand.

— Names all seven operators and roles but no usage mechanics

Your security team just got bigger (202…

THINNER COVERAGE BELOW

02 **MCP server control of Agentic Operators** Not checked yet NEW

55

All Agentic Operators can be managed through Wallarm's MCP server, letting them be plugged into external agentic workflows (triggering operators like Sift, Radar, or Wally) without returning to the Wallarm console.

Plug Wallarm Agentic Operators into an external agentic workflow via the MCP server so your own agents can trigger Sift, Radar, or Wally without opening the Wallarm console.

> 📍 Contact your customer success team to enable Agentic Operators, then connect your agentic workflow to Wallarm's MCP server to invoke operators such as Sift, Radar, and Wally programmatically.

— Names MCP server integration with a connect-and-invoke step

Your security team just got bigger (202…

03 **Autonomous mode with audit logging** Not checked yet NEW

35

Agentic Operators support an autonomous mode in which routine, approved actions run without manual confirmation, with every finding logged to the audit trail.

— Describes behaviour but no config key or toggle given

Your security team just got bigger (202…

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Next-Gen WAF agent 4.84.0 (2026-11-19) BLOG ↗

RANK

Fastly is a cloud edge computing platform that provides CDN, DDoS protection, load balancing, and real-time content delivery for websites and applications.

> Fastly's Next-Gen WAF agent 4.84.0 improves Envoy/extproc observability and strengthens auth-related threat detection.

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

01 **Envoy/extproc connection metrics and version reporting**

Not checked yet IMPROVED 40

The Next-Gen WAF agent's Envoy/extproc module now adds connection metrics and Envoy version reporting.

— Names the module but no metric fields or usage path

Next-Gen WAF agent 4.84.0 (2026-11-19)

02 **Improved INSECURE-AUTH detection** Not checked yet IMPROVED

25

The Next-Gen WAF agent now has improved INSECURE-AUTH detection.

— No detail on what changed in detection logic Next-Gen WAF agent 4.84.0 (2026-11-19)

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Sublime Security

Feedback ⓘ

Platform v1.90.262 (Verdict-based RBAC) (2026-11-20) BLOG ↗ RANK

Sublime Security is an email security platform that detects and blocks phishing, malware, and other email-based threats using customizable, human-readable detection rules.

> Sublime Security added verdict-based RBAC, letting custom roles be scoped to specific message verdicts within each flag category.

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

01 **Verdict-based role access control** Not checked yet NEW

78

Custom roles can now restrict which message verdicts — Malicious, Suspicious, Unknown, Error, Graymail, Spam, or Benign — a role can view, configured per message flag type (flagged, unflagged, user reports, DLP flagged) when creating or editing a custom role under Settings > Roles.

Limit a tier-1 analyst role so it can only open Spam and Graymail messages — keeping Malicious and Suspicious content restricted to senior responders

> 📍 In the platform, go to Settings › Roles, create or edit a custom role, and under each message flag type (Flagged, Unflagged, User Reports, DLP Flagged) select only the verdicts that role should see (e.g. Spam, Graymail) — leaving Malicious, Suspicious, Unknown, Error, and Benign unchecked.

— Names exact verdict types, flag categories, and UI path.

Platform v1.90.262 (Verdict-based RBAC)…

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ changelog · Nov 20, 2026 CHANGELOG ↗ RANK

Plerion is a cloud security platform that provides unified visibility, risk prioritization, and remediation across cloud environments to help organizations manage and reduce security risks.

> Plerion's November changelog adds six new end-of-life checks spanning AWS, GCP and Kubernetes services, tightens auto-triage of likely-exploited vulnerabilities, and fixes false-pass behavior in public-accessibility checks for RDS and EMR.

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

01 End-of-life checks for cloud data and compute services

Not checked yet NEW 80

Six new checks flag services running past end of support: `PLERION-AWS-1264` for EMR clusters on a release past end of standard support (covering release labels such as `emr-6.15.0` and `emr-spark-8.2.0` and AMI versions), `PLERION-GCP-354` for Dataproc clusters on an image version past end of support (judged by major/minor image version), `PLERION-K8S-187` for Kubernetes nodes whose kubelet minor version is past upstream end of life (skipping vendor builds like EKS, GKE, k3s, RKE2, OpenShift, AKS), `PLERION-GCP-355` for Cloud Composer environments on an unsupported image version (covering Composer 1, 2, and 3), `PLERION-AWS-1265` for Neptune clusters on an engine version past end of life (covering four-part versions such as `1.2.1.0` and `1.4.6.1`), and `PLERION-AWS-1266` for CodeBuild projects using a deprecated build image (including `standard:3.0`, `amazonlinux2-x86_64-standard:2.0`, `standard:4.0`, `amazonlinux2-aarch64-standard:1.0`, `amazonlinux2-x86_64-standard:3.0`, and all `WINDOWS_CONTAINER` environment types). All mark failed findings as End of life.

— Names every new check ID and the exact versions/images each covers.

changelog-20261120-022e84ea

02 EPSS-based severity floor in auto-triage

Not checked yet IMPROVED 70

The auto-triage rules Memory safety, no known exploit and Kernel, no known exploit now leave a vulnerability at its advisory severity when its EPSS score is above 66%, reversing any prior reduction to Medium.

— Names the exact rules and the 66% EPSS threshold. changelog-20261120-022e84ea

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Weaveworks eksctl

Feedback ⓘ v0.232.0 NOTES ↗ RANK

Eksctl is the official CLI for Amazon EKS

> eksctl v0.232.0 adds a flag to skip CloudFormation termination protection and a new cluster config field for elastic load balancing in Kubernetes network settings, alongside tighter IRSA scoping for CSI add-ons.

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

01 Elastic load balancing in Kubernetes network config

Not checked yet NEW 75

A new `elasticLoadBalancing` field under `kubernetesNetworkConfig` lets you configure elastic load balancing in cluster network settings, e.g. setting `enabled: true` in `cluster.yaml`.

Enable elastic load balancing settings in your cluster's Kubernetes network configuration.

yaml

```yaml
apiVersion: eksctl.io/v1alpha5
kind: ClusterConfig
metadata:
  name: my-cluster
  region: us-west-2
kubernetesNetworkConfig:
  elasticLoadBalancing:
    enabled: true
```

— Named config field with runnable YAML example v0.232.0

02 CloudFormation termination protection opt-out flag

Not checked yet NEW 70

The new `--cfn-disable-termination-protection` flag lets you opt out of CloudFormation stack termination protection.

— Names exact flag but no further mechanism detail v0.232.0

THINNER COVERAGE BELOW

03 Default EKS version when unset Not checked yet IMPROVED 50

When `--version` is not set, eksctl now uses the EKS default Kubernetes version instead of requiring an explicit value.

— Explains behaviour change but no command example v0.232.0

04 **Tighter IRSA scoping for CSI and EMR roles** Not checked yet

IMPROVED

50

IRSA trust for EBS and EFS CSI add-on roles now scopes to the add-on service account, reducing over-permissioning, and the EMR containers role now includes get and list secrets permissions.

— Describes permission change with named roles, no exact policy shown v0.232.0

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

OAuth authentication for webhook destinations will soon be available by default for works

BLOG ↗ RANK

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

> Databricks is rolling out OAuth 2.0 authentication for webhook notification destinations in compliance-profile workspaces and has made `ai_forecast` version 2 generally available with holiday effects, external covariates, and non-negative constraints.

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

01 **ai_forecast version 2 generally available** Not checked yet NEW

62

`ai_forecast` version 2 in Databricks SQL is now generally available, adding built-in holiday effects, external covariates, and non-negative forecast constraints, powered by a research-optimized time series foundation model for improved out-of-the-box accuracy.

·

— Names function and new capabilities but no usage example.

ai_forecast version 2 is now generally …

THINNER COVERAGE BELOW

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

UPDATED INTEGRATION NEW

53

In workspaces with the compliance security profile enabled, Databricks will soon enable OAuth 2.0 client-credential authentication for webhook destinations by default — Databricks fetches an access token from a configured OAuth token endpoint and attaches it to each notification request.

·
·
·

— Explains mechanism but no config key or endpoint named.

OAuth authentication for webhook destin…

NEW SUPPORTED PLATFORM Postgres 19 support for Lakebase

> Lakebase now supports Postgres 19 (2026-11-19) BLOG ↗ STABLE

> OAuth M2M authentication for the Salesforce, ServiceNow, and Jira connectors (2026-11-

BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

Elastic Security Labs

Feedback ⓘ

From posture to root cause: Investigate OpenShift compliance with Elastic Agent Builder

BLOG ↗ RANK

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

> Elastic Security Labs shipped a new OpenShift compliance integration that polls the OpenShift Compliance Operator APIs into Elasticsearch, builds latest-state indices and dashboards from that data, and layers an Elastic Agent Builder agent on top for plain-English posture investigation.

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

01 **OpenShift compliance data ingestion via HTTP JSON input**

Not checked yet NEW 90

Polls OpenShift Compliance Operator APIs every 60 seconds via the HTTP JSON input, shipping `compliancescans.compliance.openshift.io`, `compliancecheckresults.compliance.openshift.io`, and `complianceremediations.compliance.openshift.io` resources into Elasticsearch data streams. Every ingested event is tagged with `orchestrator.cluster.name` using the `CLUSTER_NAME` environment variable set on the Elastic Agent, keeping multi-cluster scan results separated in Elasticsearch.

Deploy Elastic Agent on an OpenShift cluster to start collecting compliance scan data into Elasticsearch, tagging every event with the cluster name so multi-cluster results stay separated.

```yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: elastic-agent
  namespace: elastic-agent-ocp
spec:
  replicas: 1
  template:
    spec:
      containers:
      - name: elastic-agent
        image: docker.elastic.co/elastic-agent/elastic-
agent:9.5.5
        env:
        - name: CLUSTER_NAME
          value: "prod-us-east-1"
        - name: FLEET_ENROLL
          value: "1"
        - name: FLEET_INSECURE
          value: "false"
```

Upload the custom OpenShift Security Posture Management integration package through the Kibana Fleet UI to register all data streams, ingest pipelines, transforms, and dashboards in one step.

> 📍 In Kibana, go to Fleet › Integrations › Upload Integration, select the OSPM zip package from the GitHub repository, then open Fleet › Agents, create a new Agent Policy, add the OSPM integration, set the three compliance endpoints and 60-second polling interval, generate an enrollment token, and deploy the manifest to your cluster.

— Names the APIs, data streams and env var, with a deploy example.

From posture to root cause: Investigate…

02 **Latest-state transforms for compliance indices** Not checked yet

NEW 68

Three Elasticsearch transforms — `scans`, `checks`, and `remedy` — continuously collapse time-series compliance data into latest-state indices named `ospm-compliance_-latest`, deduplicating check events by cluster and `metadata.name`.

— Names transforms and index pattern but gives no setup steps.

From posture to root cause: Investigate…

03 **Prebuilt Kibana dashboards for OpenShift compliance**

Not checked yet NEW 60

Adds three prebuilt Kibana dashboards — `[OpenShift Compliance] Posture Overview`, `[OpenShift Compliance] Overall Results`, and `[OpenShift Compliance] Auto Remediations` — installed automatically when the integration package is uploaded through Fleet.

— Names all three dashboards and the install trigger.

From posture to root cause: Investigate…

THINNER COVERAGE BELOW

04 **CEL custom rules for manual compliance checks** Not checked yet

NEW 50

Manual compliance rules can be evaluated with Common Expression Language (CEL) custom rules, with a scheduled Elastic Workflow annotating results back to the manual rule in the cluster.

— Mechanism described but no rule syntax or config example.

From posture to root cause: Investigate…

05 **Elastic Agent Builder agent for compliance investigation**

Not checked yet NEW 35

Adds an Elastic Agent Builder agent over the centralised compliance data so teams can query fleet-wide posture and investigate control failures in plain English.

— Described only in prose, no example queries given.

From posture to root cause: Investigate…

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ Download Cases exports directly in Tines (2026-11-19)

BLOG ↗ RANK

Tines is an apps, agents and automations, in the complete environment for enterprise AI.

> Tines changed how Cases exports are delivered, moving from email attachments to direct in-app downloads.

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

01 **In-app download for Cases exports** Not checked yet BREAKING

40

Cases exports are now downloaded directly inside Tines instead of being sent as email attachments. This replaces the previous delivery method entirely, so exported files must now be retrieved from within the app.



— Describes the change and behaviour but no exact UI path or endpoint named.

Download Cases exports directly in Tine…

└──▷ **BREAKING ON UPGRADE**

! Cases exports are no longer sent as email attachments; files must now be downloaded directly inside Tines.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Yeti

Feedback ⓘ 2.13.0 NOTES ↗

Yeti is a your Everyday Threat Intelligence

Yeti 2.13.0 adds settings for audit-log body control and URL-import safety, tightens observable type validation, and fixes API key revocation and expiry enforcement.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.13.0 https://github.com/yeti-platform/yeti.git
# already have the repo? check out this version:
$ git checkout 2.13.0
```

Control whether request bodies appear in the console and audit log file with the `audit_log_bodies` setting or `YETI_SYSTEM_AUDIT_LOG_BODIES` environment variable (default: true). Not checked yet

Restrict URL imports to public addresses only by setting `allow_internal_url_imports = false` or `YETI_SYSTEM_ALLOW_INTERNAL_URL_IMPORTS`; `/observables/import/url` now times out at 5 s connect / 30 s between bytes and returns 400 on failure. Not checked yet

Control the number of web workers in `webserver-prod` with the `YETI_WEB_WORKERS` environment variable (default: 4). Not checked yet

Observable values submitted to `POST /observables/` and `POST /observables/bulk` are now validated against their declared type; mismatched values are refused with a 400 (single) or skipped with a log entry (bulk). Not checked yet

`/graph/match` now lists an entity under every looked-up value it is linked to, returning one `[relationship, entity]` pair per link in `entities`; clients should group by the relationship's `source` or `target` field instead of de-duplicating by entity id. Not checked yet

+1 more

NEW INTEGRATION CIRCL OSINT feed (MISP format)

NEW INTEGRATION OffSeq Mirage feed

NEW INTEGRATION NetAPI Top 1M domains feed

NEW INTEGRATION phishunt.io phishing feed

NEW INTEGRATION IsMalicious on-demand analytic

NEW INTEGRATION DarkEye domain lookup on-demand analytic

NEW INTEGRATION SoChain wallet enrichment on-demand analytic

UPDATED INTEGRATION abuse.ch feeds now require authentication key

└──▷ **BREAKING ON UPGRADE**

- ! abuse.ch feeds (MalwareBazaar, ThreatFox, URLhaus) now require `YETI_ABUSECH_AUTH_KEY` or `[abusech]` in `yeti.conf`; without a key all three feeds fail.
- ! API key expiry logic is inverted on upgrade: keys with a future expiry date start working; keys whose expiry date has passed stop working.
- ! Deleted or disabled API keys sent as bearer tokens are now refused (previously they kept working).
- ! `POST /observables/` and `POST /observables/bulk` now reject values that do not match their declared type (400 for single add; skipped with a log entry in bulk); feeds or scripts that sent mistyped values will start logging skipped items.
- ! `/graph/match` `entities` now contains one `[relationship, entity]` pair per link instead of one per entity per request; clients that de-duplicate by entity id must switch to grouping by the relationship's `source` or `target` field.
- ! Event task `acts_on` patterns now fire on Python 3.11+, including all three shipped event tasks; `DatadogMetrics` will start reporting a `yeti.object` series with `type:acl`.
- ! `webserver-prod` now starts 4 workers by default (`YETI_WEB_WORKERS`) instead of one per host core.
- ! `/observables/import/url` now times out (5 s connect, 30 s between bytes) and returns 400 instead of 500 on a failed fetch.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.261119.0 NOTES ↗ RANK

OpenCTI is an open Cyber Threat Intelligence Platform

> OpenCTI 7.261119.0 ships a next-generation playbook editor, multi-threat-actor comparison on the MITRE ATT&CK matrix, dashboard-wide global filters, and a new server-level HTTP security config block, alongside configuration import/export for RBAC and email templates.

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

01 **RBAC and email template configuration presets** Not checked yet

NEW 65

RBAC configuration (groups, roles, policies) can now be exported and imported through a new Import Wizard with partial-import support, and the same configuration-preset mechanism now also covers export/import of email templates, including a new built-in User locked email template for account lock-out notifications.

— Names RBAC fields, Import Wizard, and the new email template 7.261119.0

02 **Next-generation playbook editor** Not checked yet IMPROVED

60

The playbook editor gains deterministic layout, node rewiring, an undo button, and a new node-deletion behaviour that re-attaches the downstream branch instead of removing it.

— Describes concrete editor mechanics but no exact UI path or command 7.261119.0

THINNER COVERAGE BELOW

03 **Server-level HTTP security config block** Not checked yet NEW

55

A new global HTTP switch controls CSP, rate limiting, body size, and redirects via a new server-level config block.

— Names the controls but not the exact config key 7.261119.0

04 **Global filter for dashboard widgets** Not checked yet NEW

55

A global filter can be applied across an entire dashboard to scope all widgets at once, for example by TLP marking or author, set from the dashboard editor's Filters panel.

Scope all dashboard widgets to a single knowledge domain without editing each one individually.

📍 In the dashboard editor, open the Filters panel and set a `global` filter (for example, by TLP marking or author); all widgets will inherit the constraint automatically.

— Includes a concrete navigation path and example `7.261119.0`

05 **Multi-actor MITRE ATT&CK comparison** `Not checked yet` NEW

45

The TTP Analyser now lets users compare multiple threat actors side-by-side on the MITRE ATT&CK matrix.

— Names the feature and location, no deeper mechanism given `7.261119.0`

06 **EPSS score filtering with null support** `Not checked yet` IMPROVED

45

Entities can be filtered by EPSS score, and results where the score is null are now included.

— Clear behaviour change, no config detail beyond that `7.261119.0`

07 **Inline image preview in entity Content** `Not checked yet` NEW

30

PNG and JPEG images can now be previewed inline in the `Content` section of entities.

— Bare description with no further mechanism given `7.261119.0`

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

OSSEC HIDS

Feedback ⓘ 4.5.0 NOTES ↗ RANK

OSSEC HIDS is an OSSEC is an Open Source Host-based Intrusion Detection System that performs log analysis, file integrity checking, policy monitoring, rootkit detection, real-time alerting and active response.

> OSSEC HIDS 4.5.0 expands rule matching with new agent, location, hostname and CDB field syntax, adds a post-baseline FIM create-tracking mode, and ships a PCRE2 pattern-testing CLI alongside Windows installer flags, daily log rotation, IIS log field extraction, an OpenSSL 4 build of ossec-authd, and broader notify_time handling.

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

01 Post-baseline FIM create tracking

Not checked yet NEW

85

Enable new post-baseline FIM create tracking with `<report_new_files_as_changed>`, which stores new files created after the initial baseline as `!+++` entries visible in `syscheck_control -i`.

Track files created after the initial FIM baseline as changes so they appear in syscheck_control reports, rather than being silently ignored.

xml

```
<syscheck>
  <report_new_files_as_changed>yes</report_new_files_as_changed>
</syscheck>
```

— Names config key, entry marker and inspection command, with runnable config. 4.5.0

02 Expanded rule-matching syntax for agents, paths and fields

Not checked yet NEW

78

Adds new rule-matching capabilities: expose the OSSEC agent identity as `agent_name` or `<agent_name>`, match log paths with `<location>` or `<location_pcre2>`, preserve device hostname from agent-forwarded syslog lines for `<hostname>` matching, and allow CDB `<list>` matching on decoded `system_name` via `field="system_name"`.

Write rules that alert on specific OSSEC agents by name, or scope a rule to a particular log path, without relying on IP address alone.

```xml
<rule id="100001" level="10">
  <agent_name>webserver-prod</agent_name>
  <location>/var/log/nginx/access.log</location>
  <description>Specific agent and log path match</description>
</rule>
```

— Names four new rule-matching syntax elements with a worked example. 4.5.0

03 **PCRE2 rule pattern testing CLI** Not checked yet NEW 70

Test a regex pattern against PCRE2 — with the same flags as a `<pcre2>` rule — using `ossec-regex -p`.

Verify that a PCRE2 pattern behaves as it would inside a rule before deploying it, catching flag or syntax mismatches early.

```
$ ossec-regex -p '^Failed password for (invalid user )?\S+ from (?P<srcip>\S+)'
```

— Exact runnable command given to test patterns before deploying rules. 4.5.0

04 **Daily rotation of ossec.log** Not checked yet IMPROVED 65

Rotate `ossec.log` daily into `logs/ossec/YYYY/Mon/ossec-DD.log`, then sign and compress it alongside alert logs.

— Describes new path and signing/compression but no reader action. 4.5.0

05 **ossec-authd built against OpenSSL 4** Not checked yet IMPROVED 60

Build `ossec-authd` with OpenSSL 4, enforcing TLS 1.2 via TLS_method() and using ASN.1 string accessors for certificate IP checks.

— Internal build/security change with no user-facing action. 4.5.0

THINNER COVERAGE BELOW

06 **Silent Windows installer skip flags** Not checked yet NEW 55

Skip IIS logging and integrity checking in silent Windows agent installs with `/IISLogging=no` or `/IntegrityChecking=no`.

— Named installer flags but no example invocation shown. 4.5.0

07 **notify_time scope extended to agent status checks** Not checked yet

IMPROVED

55

Honor `monitord.notify_time` for `agent_control` status and `remoted` send eligibility, not only monitord disconnect alerts.

— Names the config key and affected components but gives no example. 4.5.0

08 **IIS W3C log field extraction** Not checked yet NEW

40

Extract User-Agent, action, and srcport from default IIS W3C access logs.

— Bare list of extracted fields with no further detail. 4.5.0

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Service Gateway firmware version 3.0.37 (November 2026) (2026-11-18) BLOG ↗

RANK

Trend Vision One is a unified cybersecurity platform that provides extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

> Trend Vision One shipped FIPS mode control via new CLISH commands on the Service Gateway appliance, and full VPC enforcement for File Security's AWS Lambda functions with validated, per-region SSM-stored configuration.

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

01 VPC enforcement for FSS Lambda functions

Not checked yet NEW

95

File Security Storage stack version 4.5.4 for AWS lets you configure VPC placement for all FSS Lambda functions by providing `VpcSubnetIDs`, `VpcSecurityGroupIDs`, `RegionalVpcSubnetIDs`, and `RegionalVpcSecurityGroupIDs` together, or leaving all four empty to run outside a VPC. Settings are now validated at deploy and update time to catch misconfigured stacks before Lambda functions launch, and per-region VPC configuration is stored as SSM parameters under `/FSS/<StackSetName>/RegionalVpcConfig/` in every deployed region, removed automatically when the stack is deleted.

Retrieve the saved per-region VPC configuration from SSM after deployment to audit which subnets and security groups each region is using.

```
$ aws ssm get-parameters-by-path \
  --path "/FSS/fss-prod/RegionalVpcConfig/" \
  --recursive \
  --region us-east-1
```

— Names exact config keys, SSM path, and includes a runnable command

File Security Storage template version …

THINNER COVERAGE BELOW

02 FIPS mode CLISH commands on Service Gateway

Not checked yet NEW

45

New CLISH commands let administrators show, enable, or disable FIPS mode on the Service Gateway appliance directly from the command line; changes prompt for confirmation and trigger an automatic reboot.

— Describes behavior but no exact command syntax given

Service Gateway firmware version 3.0.37…

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl

Feedback ⓘ changelog · Nov 19, 2026 BLOG ↗

Cribl is the AI Platform for Telemetry, providing open, flexible infrastructure that helps enterprises manage, connect, and analyze their data across IT, security, and observability environments.

Privacy 3.1 raises recall by up to 13.85 and precision by up to 5.59 percentage points over 3.0 with no change in throughput or memory usage.

UPDATED DETECTION Privacy 3.1 models improve recall by up to 13.85 percentage points and precision by up to 5.59 points over Privacy 3.0, with no change in throughput or memory usage. Not checked yet

NEW SUPPORTED PLATFORM An updated benchmark for model evaluation uses outlier-seeded test cases — including sensitive values in Windows event XML, heavily escaped JSON, headerless CSV, and free-form ticket text — with roughly three quarters of events drawn from log types never seen in training. Not checked yet

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

oxo

Feedback ⓘ v2.25.2 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> oxo added a command for retrieving the logging service account when deploying it as an on-prem scanner.

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

01 **Logging service account command for on-prem scanners**

Not checked yet NEW 40

A new command retrieves the logging service account when running oxo as an on-prem scanner.

— Names the capability but not the exact command syntax v2.25.2

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ Inbound SCIM, Google Workspace Group Sync, and More

NOTES ↗

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

Log entries are more informative and managing authentication policies is easier.

Log entries are more informative. Not checked yet

NEW INTEGRATION Inbound SCIM 2.0 provisioning

NEW INTEGRATION Google Workspace group sync

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Varonis

Feedback ⓘ updated · Nov 20, 2026 NOTES ↗ RANK

Varonis connects data, AI, and threat detection into one complete security platform that allows companies to keep innovating, without increasing risk.

> Varonis added flexible scheduling for automated searches, supporting quarterly and annual recurrence patterns.

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

01 **Quarterly and annual search scheduling** Not checked yet NEW

60

Scheduled searches can now run at configurable month-based intervals — including quarterly and annual — with day-of-month or weekday-in-month execution options, eliminating manual re-runs of recurring reports.

— Describes mechanism and options but no UI path or config key snapshot-20261120

WAS THIS USEFUL? 👍 👎

MIND

Feedback ⓘ

Email DLP you can trust: Real-time blocking for Exchange and Gmail (2026-11-19)

BLOG ↗ RANK

MIND data security platform puts data loss prevention (DLP) and insider risk management (IRM) on autopilot - automatically identify, detect & prevent data leaks

> MIND introduces real-time outbound email DLP blocking for Microsoft Exchange Online, adding email as a new enforcement point alongside existing endpoint and browser policies.

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

01 **Real-time email DLP blocking for Exchange Online** Not checked yet

NEW 69

MIND now blocks outbound email in real time for Microsoft Exchange Online by routing mail through MIND for a policy decision before Exchange delivers it, requiring no agent installed on Exchange and no manual configuration in the Exchange admin center. Email policy components are now available in MIND's policy builder alongside existing endpoint and browser policies, making email a third enforcement point on the platform. Because enforcement happens at the mail-flow level, sends from phones, tablets, and personal devices receive the same DLP protection as sends from managed workstations.

— Mechanism and device scope given, but no config key or API named.

Email DLP you can trust: Real-time bloc…

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

ConductorOne

Feedback ⓘ Custom form enhancements (2026-11-19)

BLOG ↗ RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

> ConductorOne's request forms now support conditional field visibility and regex-based answer validation, letting teams build smarter, more guided access request intake.

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

01 Regex validation on request form fields

Not checked yet NEW

76

Text and text area fields on request forms can be validated against a regular expression — such as a ticket ID or cost center code pattern — with a custom error message. The web UI checks the answer as the requester types and Slack checks on submit, blocking submission until the answer matches.

Enforce a ticket ID format on a request form text field so free-text answers that don't match your ITSM pattern are blocked before submission.

> 📍 In the web UI, open a Text field on a request form, enable validation, enter a regular expression such as ^[A-Z]+-[0-9]+$, and add a custom error message like 'Enter a ticket ID in the format PROJ-1234'.

— Concrete regex example and exact UI steps, but no API/config surface named.

Custom form enhancements (2026-11-19)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Auth0 Security Agent is now in Early Access (2026-11-19) BLOG ↗

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 Security Agent EARLY ACCESS adds guided triage and remediation: it recommends Attack Protection and Tenant ACL changes and lets admins apply them after authenticating.

The Security Agent triages detected attacks and recommends Attack Protection and Tenant ACL configuration changes, letting an Auth0 tenant administrator review and apply each one after authenticating. Not checked yet

Re-run triage after applying recommendations to check the current status of an ongoing attack. Not checked yet

NEW DETECTION Credential stuffing attack detection

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Version 2026.47.1 (2026-11-18) BLOG ↗ RANK

Okta Identity Engine is a secure, scalable, and highly available authentication and user management for any app.

> Okta Identity Engine's 2026.47.1 release adds new audit logging for agent gateways, expands authentication options with on-demand identity verification and SPIFFE JWT-SVID client auth, and previews new device-related visibility and mapping features.

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

01 **SPIFFE JWT-SVID client authentication (GA)** Not checked yet NEW 35

Okta Identity Engine now supports SPIFFE JWT-SVID client authentication, now generally available.

— Names the standard but no setup detail Version 2026.47.1 (2026-11-18)

02 **Device Visibility for Windows** Not checked yet NEW 35

Okta Identity Engine adds Device Visibility for Windows, covering Desktop MFA and Device-Bound SSO, in Early Access.

— Names covered scenarios but no setup or UI path Version 2026.47.1 (2026-11-18)

03 **Agent Gateway System Log events** Not checked yet NEW 30

Okta Identity Engine now emits Agent Gateway System Log events, allowing auditing of agent gateway activity.

— Names the log event type but no schema or fields Version 2026.47.1 (2026-11-18)

04 **On-demand identity verification (Preview)** Not checked yet NEW 30

Enable on-demand identity verification as a self-service feature, currently available in Preview EARLY ACCESS.

— Describes feature and availability stage only Version 2026.47.1 (2026-11-18)

05 **Smart Card certificate policy AMR mapping** Not checked yet NEW

25

Okta Identity Engine adds Smart Card certificate policy AMR mapping, available in Early Access Preview.

— No mechanism or policy configuration detail given Version 2026.47.1 (2026-11-18)

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale Feedback ⓘ

Log streaming integration with CrowdStrike Falcon Next-Gen SIEM (2026-11-19)

BLOG ↗

Tailscale creates private WireGuard-based mesh networks with centralized identity and access controls.

NEW INTEGRATION CrowdStrike Falcon Next-Gen SIEM log streaming

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

Risk Ledger

Feedback ⓘ changelog · Nov 20, 2026 CHANGELOG ↗

Risk Ledger is a cutting-edge platform for real-time supply chain security, which enables organisations to Defend-as-One because every link matters in cybersecurity.

Risk Ledger adds an AI-assisted Domain Review Agent for all clients, PDF Review Reports on completed reviews, inline community signals in assessments, and a flexible re-review date picker.

Domain Review Agent is now on by default for all clients, running an AI first pass across all domain controls in a supplier assessment and surfacing flagged controls under an AI Findings on controls filter with a Findings (AI) tab per control. Not checked yet

Review Reports now auto-generate a PDF when a supplier review is completed, downloadable from the review drawer, containing the recommendation, approval decisions with who/what/when, non-compliant controls, open risks, remediations, full assessment answers, and optionally the AI supplier summary. Not checked yet

Community signals shared by peers now appear inline inside a supplier's assessment view, with a per-control icon and a filter to show only controls that carry a signal, for communities with review sharing switched on. Not checked yet

Re-review dates for completed reviews can now be adjusted from the supplier profile, with options from 1 month to 2 years or a custom date. Not checked yet

Decision makers now receive an email notification linking directly to the review when they are asked to decide; this can be turned off in notification settings. Not checked yet

+2 more

NEW COMPLIANCE FRAMEWORK Pre-made framework policies

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

LogicGate

Feedback ⓘ

How Do You Tell Real Agentic AI From Hype in GRC? A Buyer's Checklist (2026-11-19)

BLOG ↗ RANK

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

> LogicGate launched GRC Agents in Risk Cloud, bringing agentic AI into Third-Party Risk Management, AI Governance, Enterprise Risk Management, and Business Continuity Management applications with full audit logging and configurable triggers.

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

01 **GRC Agents for TPRM, AI Governance, ERM, BCM** Not checked yet

NEW 80

Risk Cloud now includes GRC Agents that can act within Third-Party Risk Management, AI Governance, Enterprise Risk Management, and Business Continuity Management applications. Every agent action is captured in the same Record History and Field History audit trail used for human users, including confidence score, reasoning, and source citations. Agent runs can be triggered manually or automatically based on a workflow signal, configured per application, and each agent operates under its own agent user identity governed by the same permission sets as human users, which admins can enable or disable at any time. Agent inputs and outputs are processed through OpenAI via API, retained for up to 30 days for safety monitoring, and are not used to train models.

Review what a GRC Agent changed on a vendor intake record — open the audit trail immediately after an agent run to verify every field, link, and record the agent touched.

> 📍 In Risk Cloud, open the vendor intake record the agent processed, then navigate to Record History and Field History to see each change the agent made, its confidence score, reasoning, and source citations.

— Explains mechanism, identity, logging and data handling with named apps.

How Do You Tell Real Agentic AI From Hy…

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 16th, 2026 NOTES ↗

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

Compyl adds custom fields on Controls (admin-defined, editable in drawer and grid, importable via CSV) in this release.

Admins can now define custom fields on Controls; users can view and edit them in the control drawer, see them as columns in the Controls grid, and populate them via CSV import. Not checked yet

UPDATED INTEGRATION Microsoft Defender machine inventory antivirus health

NEW INTEGRATION Google Drive publishing for Policies

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.3.0 (2026-11-19) BLOG ↗

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

Runlayer adds per-virtual-key LLM Gateway routing, dynamic subagents, an agent kill switch, audit-logged grant refusals, `tools/list` policy enforcement, `sort_by=recommended` on the agents API, and HTML artifact publishing.

Give a virtual key its own LLM Gateway routing rules and fallbacks using `runlayer_llm_gateway_usage_rule` with `scope = "virtual_key"` and `runlayer_llm_gateway_routing_rule` in Terraform, or set both in the key's Edit dialog under Settings → LLM Gateway → Virtual keys. Not checked yet

Manage MDM configurations as code with the new `runlayer_mdm_configuration` Terraform resource in the next Terraform provider release. Not checked yet

Agents on the Pi harness can now start dynamic subagents using the new Agent tool, with read-only, planning, or general-purpose workers that run in the background and use the agent's connectors, skills, and identity, with at most two running at once. Not checked yet

Sort agents by recent organizational usage with `sort_by=recommended` on the agents API and the `list_agents` MCP tool; the response also includes `users_30d`, `runs_30d`, and `last_used_at` per agent. Not checked yet

An agent harness can now request connector tools in a flat list by sending the `X-Runlayer-Tool-Exposure: force-flat` header, including on agent account endpoints and Plugins with Dynamic tools; `force-nested` keeps clients on `search_tools` and `execute_tool`. Not checked yet

+13 more

NEW INTEGRATION Google Vertex AI LLM Gateway provider

UPDATED INTEGRATION Google AI Studio and Vertex AI default Gemini models

UPDATED SUPPORTED PLATFORM Stdio connectors upgraded to Python 3.14

└─▷ BREAKING ON UPGRADE

! Stdio connectors that run inside the Runlayer backend now run on Python 3.14 and start in `/tmp` instead of `/app`; a `uv` script whose `requires-python` excludes 3.14 no longer starts. If a connector relied on `/app`, set `cwd` in its config.

! A `uv` stdio connector that runs its own project now builds a fresh environment on each start; set `UV_ISOLATED=0` in its config to reuse the project's `.venv` again.

- ! Any `--python <version>` or `-p <version>` argument must be removed from a `uv` stdio connector's args — ClickHouse connectors added from the catalog before this release include `--python 3.13` and won't start until it's removed.
- ! Sending plain text for Google Slides via the Google Workspace connector now fails with an error message instructing the caller to send a `.pptx` instead; previously it silently created a Google Doc. Existing Google Drive connectors need to be redeployed to receive file conversion.
- ! The `seen_by_user_count` API field on connector tools now stops at 50 for tools seen by 50 or more users.
- ! The deprecated connections API no longer widens an existing connection.
- ! The session filter on the Audit Logs page is now offered only to organizations with Sessions.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Red Teaming & Testing

ifixai

Feedback ⓘ v4.4.0 NOTES ↗ RANK

Ifixai is an independent Auditing of AI Agents. Run by human or the agent itself, to answer the most crucial question in the AI Agent Economy.

> ifixai's v4.4.0 release adds a full self-check CLI with SARIF and PR-comment output, trace import from five observability platforms, CI integrations for GitHub Actions/GitLab/pre-commit, and a Pro audit offer that now surfaces by default and can auto-install the iFixAi MCP server across 14 coding agents.

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

01 **Agent self-check via `ifixai audit`** Not checked yet NEW

93

`ifixai audit <agent-url>` runs the `smoke` suite and any pinned findings against an OpenAI-shaped chat endpoint to self-check an agent; pass `--yes` to confirm the sandbox prompt, set `IFIXAI_AGENT_KEY` for the agent's key, and use `--judge-provider mock` to skip external calls. Exit codes are now defined: `0` means nothing failed, `1` means a check or pinned finding failed, and `2` means no verdict was reached.

— Full command, flags, env var and exit codes given v4.4.0

02 **Pro audit offer and MCP server auto-install across coding agents**

Not checked yet BREAKING

93

`ifixai offer` prints the free Pro audit offer for a coding agent; pass `--no-promo` or set `IFIXAI_NO_PROMO=1` to skip it. `ifixai offer --add-server` adds the iFixAi MCP server (`https://mcp.ifixai.ai/mcp`) to every coding agent found on the machine, covering Claude Code, Codex, Gemini CLI, GitHub Copilot CLI, Devin, Cline, Amp, opencode, Cursor, VS Code Copilot, Kiro, Junie, Warp and Antigravity. As of this release every real `ifixai run` with a person at the terminal now opens with this offer by default; picking the Pro audit ends the run before any call to the agent and exits `0` .

— Exact command, URL, env var and full list of 14 agents named v4.4.0

03 **Trace import from five observability platforms** Not checked yet

NEW

83

`ifixai traces import` ingests trace data from OpenTelemetry GenAI, Langfuse, LangSmith, Arize/OpenInference and Braintrust exports into a trajectory capture; input over 64 MiB or 100,000 events is refused rather than truncated.

— Command, five named sources and explicit size limits v4.4.0

04 **CI integrations for self-check in GitHub Actions, GitLab CI and pre-commit**

Not checked yet NEW 83

`actions/self-check` runs the self-check in GitHub Actions, uploads SARIF to code scanning, and can post or update a verdict comment. `ci/gitlab/ifixai-self-check.gitlab-ci.yml` runs it in GitLab CI, keeping the report and SARIF as artifacts and optionally posting the verdict on the merge request. The `.pre-commit-hooks.yaml` pre-commit integration runs the self-check on push and checks the regressions file on commit.

— Three named integration files/paths with exact behavior described v4.4.0

05 **Findings explain, validate, SARIF and PR comment commands**

Not checked yet NEW 81

`ifixai findings explain` explains a self-check failure in plain language (what the agent was asked, what it should have done, and the framework controls it covers); `ifixai findings validate-regressions` checks that `.ifixai/regressions.yaml` is well-formed; `ifixai findings sarif` writes SARIF 2.1.0 output from self-check results; and `ifixai findings comment` writes a pull-request comment.

— Four commands named with exact syntax and output format v4.4.0

THINNER COVERAGE BELOW

06 **Default model updates for Gemini and Bedrock providers**

Not checked yet 45

BREAKING

The Gemini default model changed from `gemini-2.0-flash` to `gemini-3.6-flash`, and the Bedrock default model changed to `global.anthropic.claude-sonnet-4-6`.

— Names exact model identifiers but no migration guidance v4.4.0

└──▷ **BREAKING ON UPGRADE**

! Every real `ifixai run` with a person at the terminal now opens with the free Pro audit offer; pass `--no-promo` or set `IFIXAI_NO_PROMO=1` to skip it, and picking the Pro audit ends the run before any call to your agent and exits 0.

! The Gemini default model changed from `gemini-2.0-flash` to `gemini-3.6-flash`.

! The Bedrock default model changed to `global.anthropic.claude-sonnet-4-6`.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

joplin

Feedback ⓘ v3.7.24 NOTES ↗ RANK

Joplin - the privacy-focused note taking app with sync capabilities for Windows, macOS, Linux, Android and iOS.

> Joplin v3.7.24 introduces a first-class OpenAI provider built on the Responses API for AI-assisted notes, alongside a set of end-to-end encryption safeguards including key fingerprint validation and a breaking change disabling unencrypted uploads by default.

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

01 E2EE security safeguards and key validation

Not checked yet

55

BREAKING

Joplin adds a UI to validate public key fingerprints when sharing notebooks, letting users confirm the identity of share recipients; now asks for confirmation before disabling end-to-end encryption to prevent accidental deactivation; detects invalid public-private key pairs and allows recovering from them; and disables unencrypted uploads by default when E2EE is enabled, a breaking change for setups that relied on unencrypted uploads alongside E2EE.

– Describes what changed but no exact UI path or flag v3.7.24

02 OpenAI provider via Responses API

Not checked yet NEW

40

Joplin adds a first-class OpenAI provider built on the Responses API for AI-assisted note features.

– Names the Responses API but no setup steps or config key v3.7.24

└──▷ **BREAKING ON UPGRADE**

! Unencrypted uploads are now disabled by default when E2EE is enabled; existing setups that relied on unencrypted uploads alongside E2EE will need to review their sync configuration.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 20, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue