THE CYBER TOOLCHAIN

NO. 139

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 14, 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 220 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 65 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. BBOT's new BlastDNS engine cuts subdomain scans by roughly 40% and turns CDN/WAF port filtering on

vim ? help

by default, so recon spends less time on dead ports. Finite State now flags KEV membership on CVE feeds, so exploited bugs sort to the top without a second lookup. Findomain chains nuclei and ffuf straight onto discovered hosts.

IDENTIFY

Finish an external attack-surface sweep in less wall-clock time, with fewer junk hosts

BBOT swaps massdns for BlastDNS for roughly 40% faster subdomain resolution and filters CDN/WAF ports by default, which kills the familiar problem of a scan reporting every edge IP as having 80/443 open. Findomain widens passive sourcing to 55 providers and moves default storage to SQLite, so large enumeration runs are queryable rather than a pile of text files.

bbot · findomain

PROTECT

Rank CVEs by what is actually being exploited before anyone opens a ticket

KEV membership now rides on the CVE update feed and findings sort by severity across five endpoints, so the triage step of cross-referencing CISA's catalog by hand disappears from the pipeline. Dependency-sourced BOM components also carry match evidence, which is what an engineer asks for first when disputing a transitive finding.

Finite State Platform

PROTECT

Scan container registries without parking a long-lived credential next to them

Registry Scanner authenticates to Azure Container Registry via managed or workload identity, removing the static registry secret that otherwise sits in config and never rotates. Coverage extends to ROSA with hosted control planes reading Amazon ECR.

Sysdig Secure

DETECT

See tunneled and encapsulated traffic that previously walked past the sensor

Maltrail 3.6 decapsulates ERSPAN and MPLS — the formats a span port or carrier link actually delivers — and adds heuristics for ICMP tunneling and

IDN homographs, while a trail-store mmap cache drops cold reloads from 0.78s to 0.08s. eCapture now attaches process context to captured TLS traffic and writes pcapng by default, so you get HTTP reconstruction and the owning process without a conversion step.

maltrail · ecapture

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ updated · Nov 14, 2026 NOTES ↗ RANK

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

> Hacktron's latest snapshot adds SLA compliance tracking to its dashboard, no-findings report downloads for Whitebox pentests, and session-cookie login for pentest targets, alongside several integration and Slack UX fixes.

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

01 **SLA compliance ring in dashboard** NEW 60

The dashboard's PR review and Whitebox scan tabs now show an SLA compliance ring per severity, with MTTR and a Past SLA count linking to overdue findings.

— Names UI location and specific metrics shown snapshot-20261114

THINNER COVERAGE BELOW

02 **No-findings reports for Whitebox pentests** NEW 55

Completed Whitebox pentests with no confirmed findings now offer Full and Executive report downloads, covering scope, methodology, zero counts, and a clear no-findings statement.

— Names report types and content but no exact download path snapshot-20261114

03 **Session-cookie auth for pentest login agent** NEW 45

Pentest auth instructions now accept pasted session cookies, so the login agent can sign in to a target without a password.

— Clear mechanism but no exact field or config name given snapshot-20261114

04 **Correct author attribution in scan notifications** IMPROVED 45

GitLab merge request and Bitbucket pull request scans now show the author as the creator in Slack finding messages and the scan list, instead of whoever connected the integration.

— Clear before/after, no configurable surface named snapshot-20261114

05 **Connection form UX for GitLab, Azure DevOps, GHE** IMPROVED 45

GitLab, Azure DevOps, and GitHub Enterprise Server connection forms now show collapsible Setup instructions and keep any connection error until you edit a field or retry.

— Describes UI change but no config key or field named snapshot-20261114

06 **Revoked connection detection for Bitbucket** IMPROVED

40

The Bitbucket integration now flags a revoked connection and offers to reconnect it from the integration sheet.

— Names UI path but mechanism is thin snapshot-20261114

07 **Slack thread completion reaction** IMPROVED

30

Hacktron in Slack now swaps its eyes reaction for a check mark when it finishes a thread, and switches back to eyes for follow-ups.

— Cosmetic detail with no actionable surface snapshot-20261114

WAS THIS USEFUL? 👍 👎

xalgorix

Feedback ⓘ v4.6.189 NOTES ↗ RANK

Xalgorix autonomous AI pentesting agents — real-time reconnaissance, vulnerability detection, and exploitation orchestration. Go + TypeScript.

> Xalgorix hardens its file tools against credential and secret leakage by expanding path-based deny-list protections, extending that deny-list to block writes as well as reads, and fixing ripgrep-based search to behave like `grep -r`.

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

01 **Credential-file protections in file tools** NEW 85

`str_replace_editor`, `search_files`, and `list_files` now refuse to touch `xalgorix.env` files, the provider key store and the auth profile store in `Data_Dir`, their temporary copies, and `/proc/*/environ`. The `read_deny` field returned by `/api/status` now reflects these newly added deny-list entries so operators can confirm the protections are active.

Inspect the current deny-list to confirm new credential-file entries are active after upgrading.

```
$ curl -s -H 'Authorization: Bearer <token>'
http://127.0.0.1:9137/api/status | jq '.read_deny'
```

– Names exact protected paths and a verifiable API field/command v4.6.189

02 **Deny-list now blocks writes, not just reads** BREAKING 80

Paths listed in `XALGORIX_READ_DENY_LIST` now block file-tool writes as well as reads. This is a breaking change: any path previously allowed for writing via the write allow-list that also appears in `XALGORIX_READ_DENY_LIST` will now be refused on write.

– Clear before/after and exact env var, but no example v4.6.189

THINNER COVERAGE BELOW

03 **ripgrep search now matches grep -r behavior** IMPROVED 50

`search_files` using ripgrep now searches the exact path it was given and includes hidden and ignored files, matching `grep -r` behavior.

– Describes behavior change but no command or flag given v4.6.189

└──▷ **BREAKING ON UPGRADE**

! Entries added through `XALGORIX_READ_DENY_LIST` now also block file-tool writes — any path previously allowed for writing via the write allow-list that also appears in `XALGORIX_READ_DENY_LIST` will be refused on write.

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

NetSPI OCInferno

Feedback ⓘ changes since 2026-09-01 CODE ↗

OCInferno adds text-based OCI attack-path analysis, NSG security-rule enumeration, per-user DB credential enumeration, and identity propagation trust detection

└──▷ TRY IT

Find the shortest privilege-escalation path from every principal to any node with MANAGE access, using graph data already collected by OCInferno.

```
$ modules run process_og_attack_paths --graph-json
/path/to/oracle_cloud_hound.json --to-manage --summary
```

Enumerate every simple escalation path to a specific resource kind (e.g. OCIGroup) and emit results as JSON for downstream processing.

```
$ modules run process_og_attack_paths --graph-json
/path/to/oracle_cloud_hound.json --to-kind OCIGroup --all-paths --json
```

Analyze OCI privilege-escalation paths without BloodHound using the new `process_og_attack_paths` module, with flags `--graph-json`, `--to-manage`, `--to-kind`, `--all-paths`, `--max-paths`, `--summary`, and `--json` to control scope and output.

Enumerate Network Security Group security rules as part of `enum_core_network` enumeration coverage.

Enumerate per-user database credentials alongside existing identity enumeration.

Detect unrestricted identity propagation trusts — trusts with no `client_claim_values` set — via the new `TABLE_IDD_IDENTITY_PROPAGATION_TRUSTS` table and `ipt_unrestricted_client_claim` check.

Track customer secret keys per identity domain user via the new `TABLE_IDD_USER_CUSTOMER_SECRET_KEYS` table.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

findomain

Feedback ⓘ 11.0.0 NOTES ↗ RANK

Findomain is a fast subdomain enumeration tool that discovers subdomains of a target domain using certificate transparency logs and various APIs.

> Findomain 11.0.0 adds built-in nuclei and ffuf scanning against discovered hosts, expands passive source coverage to 55 providers, and switches its default storage backend to SQLite while replacing the built-in port scanner with nmap — alongside SMTP email reports, webhook alerting for nuclei findings, and a wide set of breaking changes to flags, config formats and the library API.

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

01 **Nuclei vulnerability scanning integration** NEW 90

Run nuclei templates against discovered live hosts using `--nuclei`, with `--nuclei-templates`, `--nuclei-severity`, `--nuclei-tags`, `--nuclei-exclude-templates`, `--nuclei-rate-limit`, and `--nuclei-timeout` controlling the scan, plus repeatable `--nuclei-arg` to pass arbitrary nuclei flags through.

— Complete flag set, directly runnable 11.0.0

02 **Nmap replaces built-in port scanner** BREAKING 90

Port scanning now shells out to nmap via `--pscan`, `--nmap-full`, `--nmap-syn`, `--min-rate`, `--nmap-host-timeout`, and passthrough `--nmap-arg` (or `nmap_args` in the config file); the previous built-in scanner and its `--parallel-ip-ports-scan`, `--tcp-connect-threads`, and `--tcp-connect-timeout` flags are removed.

— Names old and new flags with clear migration 11.0.0

03 **SQLite as default storage backend** BREAKING 90

Results are now stored in SQLite by default, writing to `./findomain.db` when no database flags are set; use `--sqlite <file>` or `FINDOMAIN_SQLITE_PATH` to choose a path. To keep using PostgreSQL, pass at least one of `--postgres-user`, `--postgres-password`, `--postgres-host`, `--postgres-port`, `--postgres-database`, or set `postgres_connection` in the config file or `FINDOMAIN_POSTGRES_CONNECTION` in the environment.

— Default change plus full opt-out path named 11.0.0

04 **Time-budgeted passive source queries** NEW 88

New `--source-timeout` (default 30s per request), `--source-budget` (default 300s total), and `--archive-budget` (default 20s for Wayback, CommonCrawl, UK Web

Archive, and Arquivo) flags let operators bound how long Findomain waits on passive sources.

— Full flag names and defaults given 11.0.0

05 **ffuf content discovery integration** NEW 86

Run ffuf against live hosts with `--ffuf`, `--ffuf-wordlist`, `--ffuf-threads`, `--ffuf-recursion`, and `--ffuf-recursion-depth`, plus repeatable `--ffuf-arg` for arbitrary ffuf flags.

— Complete flag set, directly runnable 11.0.0

06 **New reporting and alerting integrations** NEW 81

Send scan reports over SMTP using `--smtp-server`, `--smtp-port`, `--smtp-user`, `--smtp-password`, and repeatable `--email-to` (including new subdomains plus nuclei and ffuf findings), and route nuclei findings to a dedicated Discord or Slack channel using `smart_alerts_webhook` in the config file or `FINDOMAIN_SMART_ALERTS_WEBHOOK` in the environment, capping the total time Findomain waits on a rate-limited webhook per message using `--webhook-max-wait` or `webhook_max_wait`.

— Flags named, content of reports described Both config keys and flag named 11.0.0

07 **External subdomain import replaces built-in enumeration** 75

BREAKING

`--external-subdomains` is removed; Findomain no longer runs amass or subfinder for you — use `--import-subdomains` to bring in their output files instead.

— Old and new flag named with clear migration 11.0.0

08 **Passive source list expanded to 55 sources** NEW 73

Findomain now queries 55 passive sources (up from 14), including ahrefs, censys, shodan, fofa, intelx, leakix, netlas, and 35 others; any source can be excluded by identifier with `--exclude-sources`.

— Concrete count and named sources, exclude flag given 11.0.0

09 **ARM64 Docker image and musl binary** NEW 65

The Docker image now builds for `linux/arm64` in addition to `linux/amd64`, and a static `findomain-aarch64-musl.zip` asset is published for Alpine and other musl-based systems.

— Named platforms and exact asset filename 11.0.0

10 **Windows release asset renaming** BREAKING 65

Windows release assets are renamed to `findomain-windows.zip` and `findomain-windows-i686.zip`; downloads from the old `releases/latest/download/findomain-windows.exe.zip` URL will now return 404.

— Exact old and new filenames given `11.0.0`

---

11 **Source identifier consolidation for crtsh and threatcrowd** 65

BREAKING

`crtsh` is now a single source identifier covering both the database mirror and the JSON API, so `--exclude-sources crtsh` excludes both; the `threatcrowd` source is removed entirely.

— Named identifiers and flag interaction `11.0.0`

THINNER COVERAGE BELOW

12 **Config file format support narrowed** BREAKING 50

hjson configuration files are no longer read; supported config file extensions are now `toml`, `json`, `ini`, and `yml`.

— Formats named but no migration mechanism `11.0.0`

---

13 **`--threads` flag deprecated** DEPRECATED 40

`--threads` is deprecated; use `--lightweight-threads` instead.

— Clear replacement flag named `11.0.0`

---

14 **Complete library API rewrite** BREAKING 20

The library API changed completely — every module was renamed or replaced — with no deprecation cycle provided.

— No specific module names or migration path given `11.0.0`

---

15 **Minimum supported Rust version raised to 1.89** BREAKING 20

The minimum supported Rust version is now 1.89.

— Single version fact, no further context `11.0.0`

└──▷ BREAKING ON UPGRADE

! `--external-subdomains` is removed; Findomain no longer runs amass or subfinder for you — use `--import-subdomains` to bring in their output files instead.

- ! SQLite is now the default database backend, writing to `./findomain.db`; to keep using PostgreSQL, pass at least one of `--postgres-user`, `--postgres-password`, `--postgres-host`, `--postgres-port`, `--postgres-database`, or set `postgres_connection` in the config file or `FINDOMAIN_POSTGRES_CONNECTION` in the environment.
- ! The built-in port scanner is replaced by nmap; `--parallel-ip-ports-scan`, `--tcp-connect-threads`, and `--tcp-connect-timeout` are removed and replaced by `--nmap-full`, `--nmap-syn`, `--min-rate`, `--nmap-host-timeout`, and `--nmap-arg`.
- ! Windows release assets are renamed to `findomain-windows.zip` and `findomain-windows-i686.zip`; downloads from `releases/latest/download/findomain-windows.exe.zip` will return 404.
- ! hjson configuration files are no longer read; supported extensions are `toml`, `json`, `ini`, and `yml`.
- ! `crtsh` is now a single source identifier covering both the database mirror and the JSON API, so `--exclude-sources crtsh` excludes both.
- ! `threatcrowd` source is removed.
- ! The library API changed completely — every module was renamed or replaced; no deprecation cycle was provided.
- ! The minimum supported Rust version is now 1.89.
- ! `--threads` is deprecated; use `--lightweight-threads` instead.

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Hak5 Shark Jack Display Feedback ⓘ

back in stock · Nov 13, 2026 STORE ↗

Shark Jack Display is a jack into a network and instantly run advanced recon, exfiltration, attack and automation payloads.

Incognito Black, single (sjd-incognito-single) back in stock at `shop.hak5.org` on 13 Nov 2026.

Incognito Black, Pro Bundle (sjd-incognito-pro) back in stock at `shop.hak5.org` on 13 Nov 2026.

Cyb3rpunk Clear, single (sjd-cyb3rpunk-single) back in stock at `shop.hak5.org` on 13 Nov 2026.

Cyb3rpunk Clear, Pro Bundle (sjd-cyb3rpunk-pro) back in stock at `shop.hak5.org` on 13 Nov 2026.

WAS THIS USEFUL? 👍 👎

◆ Wireless & RF

fluxion Feedback ⓘ CVE-2026-125399 CVE ↗

Fluxion is a wireless security auditing tool that creates rogue access points to capture WPA/WPA2 passwords via social engineering attacks.

`CVE-2026-125399`: Zero-Click, Unauthenticated OS Command Injection via eval() on Attacker-Controlled ESSID in Fluxion (CVSS 8.8, affects `<= 6.32`)

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ changelog · Nov 14, 2026

CHANGELOG ↗

GitGuardian SaaS is a secrets detection platform that scans Git repositories and developer workflows to find and remediate exposed credentials and sensitive data.

GitGuardian adds a secret revocation API, richer public incident fields, new audit log events for login lockout, and expanded perimeter bulk-selection controls.

Revoke a secret via the new `POST /v1/secrets/{secret_id}/revoke` endpoint, passing provider credentials in the `credentials` body field to override workspace defaults.

Filter public incidents and check runs by `commit_sha`, `pull_request_number`, or `source_id` in addition to existing filters; each check run now also returns its repository's full name and source ID.

Public incident responses now include `company_related_verdict` (values: `related`, `unclear`, `unrelated`) and `agents_stage` in list, detail, and CSV export, and both fields are filterable.

Filter the Teams page by `Synced from` provider when Team Sync is enabled, to distinguish synced teams from manually created ones.

Filter public incidents by `Source owner` — the GitHub organization or user, or Docker Hub namespace, that published the source.

+2 more

NEW DETECTION Kubeconfig honeytoken type

UPDATED INTEGRATION JFrog Container Registry — federated Docker repository support

└─▷ **BREAKING ON UPGRADE**

! SharePoint Online and OneDrive integrations no longer require the `Organization.Read.All` permission on your Entra application; installations without it are now named after their tenant ID rather than the organization name.

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.167.10 NOTES ↗ RANK

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

> Infisical refined access request granularity for glob-based policies and loosened restrictions on self-hosted webhook alert destinations.

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

01 **Sub-path access requests within glob policies** IMPROVED

42

Users can now request access to a narrower path inside a glob-based access policy, scoping the request to a specific secret path rather than the entire glob pattern.

— Explains behaviour but no UI path or flag named v0.167.10

02 **Self-hosted webhook alerts over HTTP and internal addresses**

IMPROVED

23

Self-hosted Infisical instances can now send webhook alerts over plain HTTP and to internal addresses, which were previously restricted.

— Only mentioned in summary, no mechanism or config detail given v0.167.10

WAS THIS USEFUL? 👍 👎

Truffle Security trufflehog

Feedback ⓘ

announced · Nov 11, 2026 BLOG ↗

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

TruffleHog Enterprise adds a Connections hub, per-source detail pages, a background CSV export API (`POST /api/v3/secrets/exports` , `GET /api/v3/secrets/exports/{id}`), and a low-memory Git scanning preview mode.

Adds a Connections page that groups Sources, Notifiers, Cloud Analyzers, and Scanners into one sortable, column-configurable view.

Adds a per-source detail page showing idle/scanning status, next scheduled scan, configuration, scanner details, triage status, and recently discovered live secrets.

Export secrets to CSV at scale via `POST /api/v3/secrets/exports` and poll `GET /api/v3/secrets/exports/{id}` for status; filters apply via query string.

Adds a low-memory Git scanning mode for scanners on memory-limited hosts (private preview).

NEW DETECTION New detectors: SonarQube, Lambda Cloud, Fireworks AI, and xAI analyzer

└─▷ **BREAKING ON UPGRADE**

! The Integrations page is renamed to Connections across navigation and the product; existing links to the old Scanners page under Settings redirect to the new Scanners tab under Connections.

WAS THIS USEFUL? 👍 👎

◆ Code Security

GitHub Feedback ⓘ

GitHub Enterprise Server support in the GitHub Copilot app (2026-11-13) BLOG ↗

RANK

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

> GitHub Copilot gained GitHub Enterprise Server account support and a new Gemini 4 Hydrogen model option, while the REST API added repository-level controls for AI Scan.

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

01 **AI Scan enablement via REST API** NEW 95

New `repository_scan` property can be set to `enabled` or `disabled` on the AI Scan enablement endpoints, alongside `pr_scan`, in the same request via `PATCH` to `/orgs/{org}/code-scanning/ai-scan` or `/repos/{owner}/{repo}/code-scanning/ai-scan`; if either change is blocked a `422` error is returned and neither setting changes. Organization-level `repository_scan` cascades to repositories: an org-level disable prevents repos from re-enabling it, an org-level enable lets repos inherit and individually opt out, and enterprise policies that disallow AI Scan cause enable requests to return `422`.

Enable AI Scan for repositories across an entire organization while also enabling PR scan in one API call.

```
$ curl -X PATCH \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  https://api.github.com/orgs/{org}/code-scanning/ai-scan \
  -d '{"repository_scan": "enabled", "pr_scan": "enabled"}'
```

Opt a single repository out of AI Scan for repositories when the org has it enabled but you want to exclude a specific repo.

```
$ curl -X PATCH \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/vnd.github+json' \
  https://api.github.com/repos/{owner}/{repo}/code-scanning/ai-scan \
  -d '{"repository_scan": "disabled"}'
```

— Exact endpoints, methods, fields, and cascade behaviour with error codes

Manage AI Scan for repositories using t…

02 **Gemini 4 Hydrogen model in Copilot** NEW 75

GitHub Copilot now supports Gemini 4 Hydrogen, Google's latest Flash model, selectable from the model picker in Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, the Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, Eclipse, and the Copilot SDK. Copilot Enterprise and Business administrators can manage access via the model policy in Copilot settings; under default model enablement, the model is automatically enabled unless the global default is off or the model is explicitly disabled.

Select Gemini 4 Hydrogen for a Copilot CLI session to run fast, multistep agentic coding tasks from the terminal.

> 📍 In the Copilot CLI model picker, select Gemini 4 Hydrogen before issuing your request.

Restrict Gemini 4 Hydrogen access org-wide before its gradual rollout reaches all users.

> 📍 In Copilot settings (Enterprise or Business admin), go to the model policy and explicitly disable Gemini 4 Hydrogen under Default model enablement.

— Lists all surfaces and admin controls, lacks performance specifics

Gemini 4 Hydrogen in GitHub Copilot (20…

03 **GitHub Enterprise Server accounts in Copilot app** NEW 72

The GitHub Copilot app now generally supports connecting a GitHub Enterprise Server account to work with remote GHES repositories, issues, and pull requests, available in app version 1.2.0 and later. Add it via Settings > Accounts > Add account, then choose GitHub Enterprise Server, entering the GHES instance URL and credentials.

Add your GHES account to the Copilot app to access on-premises repositories, issues, and pull requests alongside your `GitHub.com` work.

> 📍 1. Sign in with a `GitHub.com` or GHE.com account. 2. Go to Settings › Accounts › Add account. 3. Choose GitHub Enterprise Server and enter your GHES instance URL and credentials.

— Names exact UI path and version but limited mechanism detail

GitHub Enterprise Server support in the…

WAS THIS USEFUL? 👍 👎

syzkaller

Feedback ⓘ changes since 2026-11-12 CODE ↗

Syzkaller is an unsupervised coverage-guided kernel fuzzer developed by Google that automatically finds bugs in OS kernels by generating and executing random syscall sequences.

syzkaller adds cross-compilation and image-creation support for FreeBSD and OpenBSD on Linux, plus a `/healthz` endpoint and kernel_builder_image option for syz-ci

└──▷ **TRY IT**

Build an OpenBSD/amd64 target image, sysroot, and builder image on a Linux host for cross-compiled fuzzing.

```
$ tools/create-openbsd-image.sh -o ./openbsd-image -v snapshots
```

Check whether a syz-hub instance is alive in a health-monitoring or deployment pipeline.

```
$ curl http://<syz-hub-host>:<port>/healthz
```

Configure a custom kernel builder container image in syz-ci with the new `kernel_builder_image` option.

Add a `/healthz` endpoint to `syz-hub` for liveness checks and graceful shutdown support.

Cross-compile and build FreeBSD kernels on Linux hosts, using the new `tools/create-freebsd-image.sh` script to produce target disk images.

Cross-compile and build OpenBSD/amd64 kernels on Linux hosts, using the new `tools/create-openbsd-image.sh` script to produce target disk images, sysroots, and builder images.

`syz-hub` now validates manager names on registration.

WAS THIS USEFUL? 👍 👎

Snyk Code

Feedback ⓘ changelog · Nov 13, 2026 BLOG ↗

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

Snyk Code now reads project-own configuration to resolve imports in Go, TypeScript, and JavaScript, following data across packages and modules for deeper analysis.

NEW DETECTION Kong Gateway plugin analysis

NEW DETECTION New Python detections: Django CSRF, Flask response handling, JWT signature verification

NEW DETECTION New JavaScript detections: command injection, prototype pollution

NEW DETECTION New PHP detection: cookie security settings

NEW DETECTION New Java detection: SQL injection

UPDATED SUPPORTED PLATFORM PHP up to 8.5 and Swift up to 6.4 support

UPDATED SUPPORTED PLATFORM JavaScript and TypeScript support for `using` declarations

WAS THIS USEFUL? 👍 👎

Checkmarx

Feedback ⓘ

The Security Backlog Can’t Keep Up with AI-Speed Development (2026-11-13)

BLOG ↗

Checkmarx agentic application security software platform Combines Hybrid scanning, AI-powered agents, and unified risk intelligence across every attack surface

Checkmarx adds an OpenAI Daybreak Defense Factory integration, piping SAST findings via the Codex Security CLI into AI-powered investigation and remediation workflows.

Import Checkmarx SAST findings into OpenAI Daybreak Defense Factory by running the Codex Security CLI with `--provider checkmarx` and `--dry-run` to preview, then without `--dry-run` to import.

NEW INTEGRATION Supported scan result formats for the Daybreak Defense Factory integration are Checkmarx One JSON and CxSAST XML reports, covering both Checkmarx One and self-managed CxSAST deployments.

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ announced · Nov 12, 2026

BLOG ↗ RANK

SonarQube is an enforce consistent quality gates across every pull request and pipeline, from IDE to CI/CD, with automated code review for bugs and security issues.

> SonarQube added an integration with OpenAI's Daybreak Defense Factory, letting Codex Security import SonarQube findings and validate fixes against the Sonar quality gate.

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

01 **Codex Security integration via Daybreak Defense Factory** NEW

50

SonarQube now works within OpenAI Daybreak Defense Factory through a SonarQube importer in Codex Security, a SonarQube plugin for Codex, and the SonarQube CLI, enabling security findings to flow into Codex Security in real time for validation and remediation against the Sonar quality gate.

— Names the plugin, importer and CLI but no commands or config shown.

launch-20261112-13f7ec48

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ announced · Nov 12, 2026 BLOG ↗ RANK

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

> Semgrep shipped a beta MCP server that lets AI coding agents triage findings, manage scans, and configure policy directly against the Semgrep AppSec Platform, alongside a wave of Guardian updates adding Copilot support, org-wide policy enforcement, AI-powered vulnerability detection, lockfile protection against supply chain attacks, and a staged-rollout Observe Mode.

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

01 **MCP server for AI coding agents** NEW INTEGRATION NEW

94

Connect Claude Code, Codex, Cursor, VS Code, or ChatGPT as an MCP client to Semgrep AppSec Platform. Agents can search and filter findings from Semgrep Code, Supply Chain, Secrets, and Agentic Workflows, triage and reopen findings in bulk, and read code snippets and finding history; start managed scans, find failed scans, read scan logs and errors, and retry scans; look up advisories, list dependencies, and find every repository affected by an advisory; dry-run and apply policy changes, manage exceptions, and configure Semgrep Guardian rules; pull reports by severity, product, or rule including median open age, findings activity, and guardrails adoption; trigger and schedule Agentic Workflows (admins only), read job status and task logs, and track credit usage; and create and link Jira and Linear tickets, build automations, view Slack notification settings, and check and sync SCM connections. Authenticate via OAuth (browser sign-in scoped to an organization) or an API token — a Web API token allows all read and write operations, a Read-only token allows reads only, and an Agent (CI) token covers read and write operations on scans, while Member tokens are not supported. A `discover` tool finds the right MCP tool for a task and a `semgrep_account` tool returns the connected organization. Available to organizations on Teams or Enterprise plans (trials included) using any MCP client that supports remote servers over HTTP.

Connect Cursor (or any MCP client that supports remote HTTP servers) to the Semgrep MCP server so your agent can triage findings and manage policies without leaving the editor.

📍 In Cursor, open Settings › MCP and add the Semgrep MCP server URL. Choose OAuth to sign in through your browser and scope access to your organization, or supply a Web API token for full read and write operations. Once connected, your agent can call the `discover` tool to list available capabilities or `semgrep_account` to confirm which organization it is connected to.

— Extensive named tools, token types and client list with a connection walkthrough.

launch-20261112-f2d6b2d0

02 **Guardian installation across coding agents** NEW

83

Guardian can now be installed in Claude Code via `claude plugin install semgrep@claude-plugins-official`, in Codex via `codex plugin marketplace add https://github.com/semgrep/guardian --sparse .claude-plugin --sparse plugin` followed by `codex plugin add semgrep@semgrep-marketplace`, and in Copilot via `copilot plugin marketplace add https://github.com/semgrep/guardian` followed by `copilot plugin install guardian-copilot@semgrep-marketplace`, extending Guardian beyond the previously supported Codex and Claude Code agents to Microsoft Copilot.

Bootstrap Guardian in Claude Code so every future session enforces your org's Semgrep policies automatically.

```
$ claude plugin install semgrep@claude-plugins-official
```

Add Guardian to Codex so the agent can detect, fix, and report Semgrep findings without leaving the agent environment.

```
$ codex plugin marketplace add
https://github.com/semgrep/guardian --sparse .claude-plugin --
sparse plugin
codex plugin add semgrep@semgrep-marketplace
```

— Exact install commands for every supported agent. launch-20261112-c889153f

03 **Org-wide policy enforcement in Guardian** NEW 72

Guardian's org-wide policy system lets security teams define custom YAML rules (such as a `no-legacy-auth` rule) in a local rules file or a repo-wide file, automatically applying them to every developer's agent session without manual distribution.

Define a custom org rule that blocks a deprecated authentication library, then commit it to your local rules file so Guardian enforces it in every agent session.

```yaml
rules:
  - id: no-legacy-auth
    languages: [python]
    severity: ERROR
    message: Use auth.v2.verify() — legacy_auth is deprecated.
    pattern: legacy_auth(...)
```

— Shows rule syntax but not exact config file path. launch-20261112-c889153f

04 **Lockfile protection against supply chain attacks** NEW 62

Guardian blocks agent-initiated changes to lockfiles that would introduce malware or a reachable vulnerability, preventing unsafe dependency changes before they reach the build system, and pushes advisories to every developer to check whether their installed version is still vulnerable when a critical supply chain incident (such as PHP2Shell or Node2Shell) is disclosed.

— Named incidents and mechanism, but it's automatic with no reader action.

launch-20261112-c889153f

THINNER COVERAGE BELOW

05 **AI-powered vulnerability detection in Guardian** NEW DETECTION

NEW

51

Guardian's AI-powered detection reviews code at generation time and reasons about complex vulnerabilities including insecure direct object references (IDOR) and broken object-level authorization (BOLA), as well as code written with malicious intent.

— Describes behavior but no configuration or trigger for readers.

launch-20261112-c889153f

NEW SUPPORTED PLATFORM Mandoline language and framework support

NEW COMPLIANCE FRAMEWORK Mandoline OWASP Benchmark and NIST Juliet test suite coverage

NEW SUPPORTED PLATFORM EU multi-tenant region

NEW INTEGRATION OpenAI Daybreak Defense Factory import integration

> announced · Nov 12, 2026 BLOG ↗
>
> Semgrep adds Mandoline, an agent-oriented static-analysis engine with CLI, MCP, and gRPC interfaces, callgraph/dataflow/taint primitives, and UniFFI bindings for six languages.
>
> Access Mandoline via CLI, MCP, or gRPC interfaces, with callgraph output available in text, JSON, GRAPHVIZ, and protobuf formats.
>
> Semgrep adds Mandoline, a decomposed static-analysis engine built for agentic workflows, exposing callgraph, dataflow, taint paths, API endpoint detection, and dependency reachability as composable primitives.
>
> Mandoline adds native library bindings for six languages via UniFFI, plus WASM support, enabling use in agent containers, edge compute, and the browser or IDE.
>
> Semgrep customers now have access to Mandoline inside Semgrep Agentic Workflows and Semgrep's Context Layer.

> announced · Nov 12, 2026 BLOG ↗
>
> Semgrep adds Mandoline (a millisecond-speed program analysis engine), Context Layer (private beta, persistent codebase context), Agentic Workflows, and Semgrep MCP for agent-native AppSec.
>
> Semgrep adds Mandoline, a next-generation program analysis engine that answers questions about code flows and reachable paths in milliseconds, scoring 100% on the OWASP Benchmark for Java and 98.28% on Python in a single run with no AI token usage.
>
> Semgrep adds Context Layer, now in private beta, giving agentic workflows persistent, dynamic understanding of a codebase and organization across runs,

enabling agents to surface 2x–7x more real vulnerabilities compared to agents without it.

Semgrep adds Agentic Workflows, powered by Mandoline and Context Layer, to automate complex detection, triage, and fix tasks following guided paths set by Semgrep's security experts.

Access the full Semgrep platform from any agent by pointing it at the Semgrep MCP server, available in Codex, Claude Code, Cursor, and CLI.

Semgrep adds deployment options including bring-your-own-key (BYOK) models, Agentic Workflows in CI, Semgrep Dedicated single-tenant hosting, and an EU multi-tenant region.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Cloudsmith

Feedback ⓘ

CLI v1.30.1 and v1.31.0: AI skills registry, Maven, Gradle, and sbt credential helpers, a

BLOG ↗ RANK

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

> Cloudsmith's CLI picked up a new AI skills registry command surface, credential helpers for Maven/Gradle/sbt, and quality-of-life fixes for deprecated option warnings and SAML workspace persistence.

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

01 **Maven, Gradle, and sbt credential helpers** NEW 50

The Cloudsmith CLI now sets up Maven, Gradle, and sbt to fetch credentials directly from the CLI, so build files no longer need a stored API key.

– Names three build tools but no exact command CLI v1.30.1 and v1.31.0: AI skills regi…

02 **AI skills registry management in CLI** NEW 30

The Cloudsmith CLI now manages AI skills from the terminal, adding a new command surface for the AI skills registry.

– No command names or flags given CLI v1.30.1 and v1.31.0: AI skills regi…

03 **Workspace persistence after SAML login** IMPROVED 30

The Cloudsmith CLI now remembers your workspace after a SAML login, removing the need to re-enter it on subsequent commands.

– Clear behaviour change but no command specifics

CLI v1.30.1 and v1.31.0: AI skills regi…

04 **Warnings for renamed CLI options** IMPROVED 25

The Cloudsmith CLI now warns when you use a deprecated option name, flagging renamed options at invocation time.

– No specific option names given CLI v1.30.1 and v1.31.0: AI skills regi…

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback

New Control: Fork Pull Request Workflows Should Require Approval (2026-11-13)

BLOG ↗

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

StepSecurity adds a Critical-severity dashboard control that flags weak fork PR approval policies at both org and repo level, with hourly refresh.

Each finding distinguishes whether the weak policy comes from the organization's default or a repository-level override, so repositories that deviate from the org setting are caught separately.

NEW DETECTION Findings update on the next hourly scan after the setting changes; to remediate, set the policy to `Require approval for all external contributors` under Settings > Actions > General in GitHub.

WAS THIS USEFUL?

◆ Web Application Security

PortSwigger Burp Suite

Feedback ⓘ

Professional / Community 2026.11 (2026-11-13) BLOG ↗ RANK

Burp Suite is a we enable the world to secure the web. Burp Suite, the Web Security Academy, and world-leading research from PortSwigger â the AppSec communityâs home.

> Burp Suite 2026.11 adds a new Academy tab for launching Web Security Academy labs directly from the tool, alongside easier project switching and restarting.

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

01 **Academy tab for launching labs in-tool** NEW 50

A new Academy tab lets users access and launch Web Security Academy labs directly inside Burp Suite without switching to a browser. Users click the Academy tab in the top navigation bar to browse and launch labs.

Access Web Security Academy labs directly inside Burp Suite without switching to a browser.

> 📍 In Burp Suite, click the Academy tab in the top navigation bar to browse and launch Web Security Academy labs.

– Names UI location and steps, lacks deeper mechanism

Professional / Community 2026.11 (2026-…

02 **Easier project switching and restart** IMPROVED 20

Burp Suite now makes switching projects and restarting the application easier from within the UI.

– Vague description with no specifics on mechanism or UI path

Professional / Community 2026.11 (2026-…

WAS THIS USEFUL? 👍 👎

Invicti

Feedback ⓘ Release 20261112 (2026-11-12) BLOG ↗

Official documentation for Invicti AppSec Platform, Invicti Platform, and Invicti Enterprise — web application security testing, DAST, API security, and vulnerability management.

Invicti adds target type filtering on the Scans page, configurable default Accept headers per target, removed-operation tracking in the API catalog, and API spec discovery in Source Scan.

Filter and sort scans by target type (Web or API) using the new target type filter and column on the Scans page.

Set the default Accept header the scanner sends on web and API requests using the `Default Web Accept Header` and `Default API Accept Header` settings in your target's advanced settings.

Source Scan now finds existing API specs in your source code and uses them to build a more complete picture of the API your application exposes.

API catalog now marks operations that no longer appear in a re-imported API definition as `Removed`, with a status filter to show or hide them.

Zero Configuration discovery is now enabled by default, requiring no extra setup.

+2 more

NEW DETECTION Sensitive data detection in API responses

NEW SUPPORTED PLATFORM GraphQL schema support in API catalog

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ JavaScript SDK 4.3.0 (2026-11-13) BLOG ↗ RANK

Fastly is a cloud edge platform providing CDN, DDoS protection, load balancing, and real-time content delivery to accelerate and secure web applications.

> Fastly's JavaScript SDK 4.3.0 adds new cryptographic key import formats, a mid-request handoff capability for compute requests, and stabilizes previously experimental APIs into standard modules.

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

01 **SPKI and PKCS8 key import formats** NEW 65

Fastly's JavaScript SDK now supports `spki` and `pkcs8` formats in `SubtleCrypto.importKey` for importing public and private keys.

— Names exact API and formats but no usage example JavaScript SDK 4.3.0 (2026-11-13)

THINNER COVERAGE BELOW

02 **Request handoff for compute requests** NEW 30

Fastly's JavaScript SDK adds a request handoff capability, allowing compute requests to be handed off mid-processing.

— No API name, flag, or mechanism detail given JavaScript SDK 4.3.0 (2026-11-13)

03 **Stabilized experimental APIs** IMPROVED 20

Formerly experimental APIs are now stabilized and documented in appropriate modules, replacing the experimental module.

— No specific APIs or modules named JavaScript SDK 4.3.0 (2026-11-13)

NEW DETECTION Best-practice workspace configurations: attack signal blocking, malformed payload rules, known malicious IP rules

> Apply recommended configurations when creating workspaces (2026-11-12) BLOG ↗
>
> STABLE
>
> Fastly now lets you apply best-practice configurations when creating a workspace in the control panel.
>
> Apply a set of best-practice configurations at workspace creation time, including immediate blocking of requests tagged with attack signals and rules covering malformed payloads and known malicious IP addresses.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Request Tracker / RTIR (Best Practical)

Feedback ⓘ Redesigned Attachments and Links in RT 6.0.4 (2026-11-13)

BLOG ↗ RANK

Request Tracker / RTIR (Best Practical) is an open source ticketing system from Best Practical. Self-host RT or let us run it in the cloud.

> Request Tracker 6.0.4 reworked how tickets handle attachments and links, adding direct uploads with pinning and bulk download, plus a redesigned Links section with type-aware autocomplete and inline field editing.

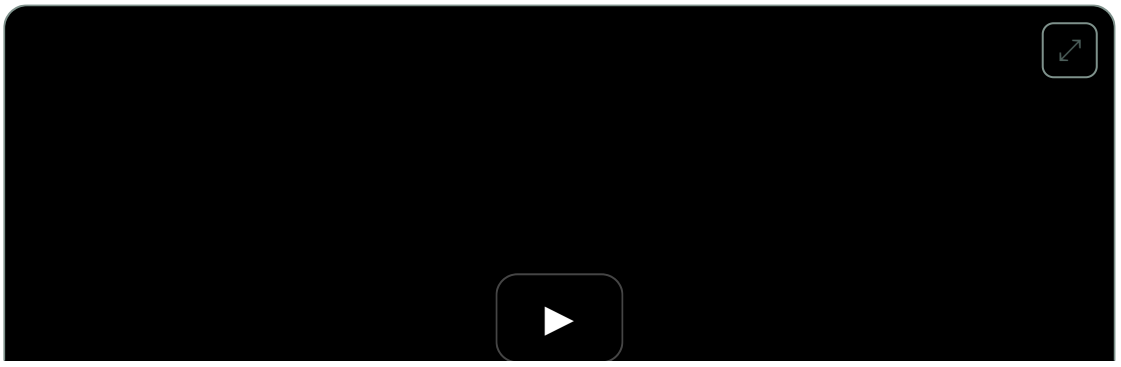
└──▷ **WHAT SHIPPED · 2 FEATURES** most completely described first ? what's the number?

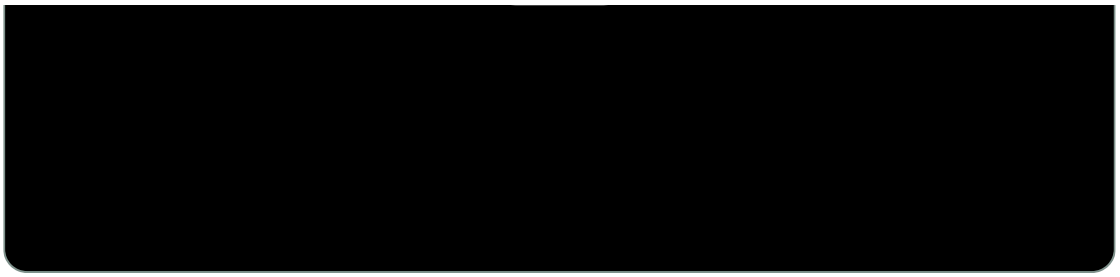
01 Direct attachment uploads with pinning and bulk download

IMPROVED 70

Attachments can now be uploaded directly from the ticket display page without starting a reply or comment, and can be renamed, pinned, or deleted from that same page. Pinning keeps key files at the top of the attachment list, and the list now supports multiple sort options, filtering, and bulk actions including bulk download.







Pin the current contract to the top of a ticket's attachment list so the whole team can find it without scrolling.

> 📍 On the ticket display page, go to the Attachments section, upload or locate the contract file, then select `Pin` from its action menu to keep it at the top of the list.

Bulk-download all attachments from a ticket to archive a closed project's files in one action.

> 📍 On the ticket display page, go to the Attachments section, select all files using the bulk-select control, then choose `Bulk download` from the bulk actions menu.

— Describes UI mechanism and concrete actions but no config keys or limits

Redesigned Attachments and Links in RT …

---

02 **Redesigned Links section with autocomplete and inline edit**

IMPROVED

70

The Links section now shows a dropdown of all linkable RT object types — assets, articles, and other objects — plus an option to add a plain web link. After choosing a link type, the input field becomes an autocomplete search box, and users can choose which fields appear for linked tickets and edit those fields inline without leaving the current ticket.

Link an asset to a ticket and update one of its fields inline without opening the asset's own page.

> 📍 On the ticket display page, go to the Links section, open the link-type dropdown and select the asset object type, search for the asset by name in the autocomplete box, add the link, then click the desired field on the linked asset row to edit it inline.

— Explains behaviour change and workflow but no API or flag names

Redesigned Attachments and Links in RT …

WAS THIS USEFUL? 👍 👎

Belkasoft

Feedback ⓘ announced · Nov 11, 2026 BLOG ↗ RANK

Accelerate your digital forensics and incident response investigations with Belkasoft

> Belkasoft X 2.13 adds MFT-based deleted file recovery alongside a wave of BelkaGPT upgrades — document summarization, email/SMS topic detection, faster Hub- and CPU-based answer generation with chain-of-thought display — plus browser-based Google Cloud sign-in, dump.bin support for MTK acquisition, and new mobile app artifact coverage.

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

01 BelkaGPT document summarization NEW

93

BelkaGPT now summarizes documents across word processor files, PDFs, spreadsheets, slide decks, and plain-text logs via Summarize with BelkaGPT in the Artifacts window's Documents profile, with Quick summary covering the first 500 words and Full summary covering up to 25,000 words, in a selectable language, with offload support via BelkaGPT Hub.

Quickly assess a large batch of documents for investigative relevance without reading each one in full.

📍 Select all items under the Documents profile in the Artifacts window, right-click and choose Summarize with BelkaGPT, select Full summary (up to 25,000 words) or Quick summary (first 500 words), pick the output language, and optionally route processing to BelkaGPT Hub if the local workstation lacks a compatible GPU.

— Named modes, word limits, and exact UI steps launch-20261111-ae8b2d91

02 BelkaGPT topic detection for emails and SMS IMPROVED

86

BelkaGPT topic detection, previously available only for chats, now also covers emails and SMS messages, using the same predefined detectors (money, drugs, operational coordination, account access) and custom detectors already available for chats, run from the BelkaGPT Topic Detection panel.

Surface relevant email and SMS threads matching investigative topics — such as drug references or operational coordination — across all communication types at once.

📍 In the BelkaGPT panel, open Topic Detection, select Emails and SMS as sources alongside any existing chat sources, choose a predefined detector (for example, Drugs or Operational coordination) or add a custom detector describing case-specific coded language, then run detection to collect matching messages.

— Named detectors and run path, scope expansion clear launch-20261111-ae8b2d91

03 **MFT-based deleted file analysis** NEW 85

Belkasoft X 2.13 adds MFT-based deleted file analysis for NTFS volumes, surfaced in the File System window under a new Deleted items node with Orphans (unknown parent path) and Path restored (reconstructed full path) subnodes, filterable by name, size, or timestamp, with recovery via Save.

Locate and recover deleted files from an NTFS volume without leaving the case — useful when a suspect wiped files but data may still be on disk.

📍 In the File System window, expand the target NTFS volume, then open the Deleted items node. Browse Orphans for files whose parent path is unknown or Path restored for files whose full original path was reconstructed. Filter by name, size, or timestamp columns to narrow results, then right-click a file and choose Save to recover it.

— Exact navigation and filtering steps given launch-20261111-ae8b2d91

04 **BelkaGPT answer generation: Hub offload, fast CPU mode, chain of thought**

IMPROVED 70

BelkaGPT answers can now be generated via BelkaGPT Hub to offload Q&A processing from the local forensic workstation, and a new fast answer mode on CPU reduces answer time to about 1 to 3 minutes (down from 12 to 15 minutes) on machines without a compatible GPU or Hub. BelkaGPT also now shows chain of thought during answer generation, displaying each reasoning step and the time it took.

— Concrete timings given but no exact invocation path launch-20261111-ae8b2d91

05 **MTK acquisition dump.bin support and decryption utility** NEW 70

MTK acquisition now supports `dump.bin` files added directly to a case as a mobile image, as an alternative to `.belkaml` images. A new utility, available on request from support@belkasoft.com, creates decryption material for MTKClient dumps so they can be decrypted in Belkasoft X.

— Names file type and contact point for utility launch-20261111-ae8b2d91

06 **New Android and iOS artifact support** NEW 60

Belkasoft X adds Android artifact support for `Claude AI v1.260625.20`, `Copilot v30.0.440623001`, `ChatGPT v5.2`, `Google Files v1.2263`, Google Messages, Google Photos, `Samsung Gallery Trash v14.5.12`, and `Samsung My Files`

`v14.5`, plus iOS artifact support for `Claude AI v1.260618.1` and `Copilot v30.0.440608001`.

— Exhaustive named app versions but no usage steps launch-20261111-ae8b2d91

THINNER COVERAGE BELOW

07 **BelkaGPT classification results in Evidence Reader exports**

IMPROVED 55

BelkaGPT classification results — topic detection, SQLite database classification, and picture classification — are now included in exports from the Artifacts window to Belkasoft Evidence Reader.

— Names export surface but no further mechanism launch-20261111-ae8b2d91

08 **Browser-based sign-in for Google Cloud acquisition** NEW

50

Google Cloud acquisition now supports browser-based sign-in with two-factor authentication for all supported Google services.

— Clear capability but no workflow detail launch-20261111-ae8b2d91

09 **Reports split by contact or conversation** IMPROVED 45

Reports can now be split by contact or conversation for Mobile Apps and SMS, including RSMF reports.

— Names report types but no mechanism detail launch-20261111-ae8b2d91

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-11-11 CODE ↗

CAPEv2 is an open-source malware sandbox that automatically executes and analyzes suspicious files to extract configurations, payloads, and behavioral indicators.

CAPEv2 adds tshark-based network processing with TLS session key parsing from tlsdump.log and NSS keylog files

CAPEv2 now parses TLS session keys via `parse_tlsdump_line` in `utils/tls.py`, supporting NSS keylog labels (`CLIENT_RANDOM` , `CLIENT_EARLY_TRAFFIC_SECRET` , `CLIENT_HANDSHAKE_TRAFFIC_SECRET` , `SERVER_HANDSHAKE_TRAFFIC_SECRET` , `CLIENT_TRAFFIC_SECRET_0` , `SERVER_TRAFFIC_SECRET_0` , `EARLY_EXPORTER_SECRET` , `EXPORTER_SECRET`) alongside the existing Cape-native `tlsdump.log` format, and filters out all-zero secrets automatically.

CAPEv2 adds tshark integration into `modules/processing/network.py` for network traffic processing, pulling in subprocess-based dissection and deduplication helpers `deduplicate_file_to_storage` and `get_files_storage_path` .

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

November 12, 2026 - Archived project monitoring, findings sort, CVE metadata, and depend

BLOG ↗ RANK

Finite State Platform is a software supply chain security tool that analyzes firmware and embedded software for vulnerabilities, risks, and compliance issues.

> Finite State Platform's latest release adds KEV membership flags on CVE update feeds, severity-ranked sorting for findings across five endpoints, match evidence for dependency-sourced BOM components, and a breaking change to how CWE identifiers are formatted in CVE metadata responses.

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

01 **KEV membership flags on CVE updates** NEW 90

New optional `kev` and `vcKev` boolean fields appear on `oldValue` and `newValue` in items returned by `GET /v0/cves/updates`, letting callers filter for CVE updates where the CVE was newly added to a Known Exploited Vulnerability catalog; such updates are served under `updateType=exploit_update`.

Poll for CVE updates and check whether a change also reflects a new Known Exploited Vulnerability catalog addition, without a separate lookup.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/api/public/v0/cves/updates' | jq
'[.items[] | select(.newValue.kev == true or .newValue.vcKev ==
true) | {cveId, updateType, kev: .newValue.kev, vcKev:
.newValue.vcKev}]'
```

— Names fields, endpoint, and updateType value with a runnable query

November 12, 2026 - Archived project mo…

02 **Severity-ranked sorting for findings endpoints** IMPROVED 90

Version-scoped findings can now be sorted by severity rank (Critical, High, Medium, Low, Info, None, Unknown) and by component group and name via `GET /v0/findings`, `GET /v0/projects/{id}/findings` (up to 10 versions), `GET /v0/scans/{id}/findings`, `GET /v0/components/{id}/findings`, and `GET /v0/versions/{id}/findings`.

Retrieve version-scoped findings already sorted by severity rank to feed a triage dashboard without post-processing.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/api/public/v0/versions/<versionI
d>/findings'
```

— Lists all five endpoints and sort criteria with a runnable example

November 12, 2026 - Archived project mo…

---

03 **CWE identifiers returned as formatted strings** BREAKING

90

`GET /api/public/v0/cves/{cveId}/metadata` and `GET /api/public/v0/findings/{projectVersionId}/{findingId}/cves` now return every `weaknesses[].description[].value` as a `CWE-<number>` string instead of a bare number (e.g. `399` becomes `'CWE-399'`); placeholder values `CWE-0` and `CWE--1` are omitted, so a CVE whose only weakness was a placeholder now returns an empty `weaknesses` list.

Confirm a CVE weakness is now returned as a `CWE-<number>` string so downstream parsers that expected a bare integer can be updated.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/api/public/v0/cves/<cveId>/metad
ata' | jq '[.weaknesses[].description[].value]'
```

— Gives exact before/after format and endpoints affected

November 12, 2026 - Archived project mo…

THINNER COVERAGE BELOW

---

04 **Match evidence for dependency-sourced BOM components** NEW

30

Finite State Platform now shows match evidence, guidance, and dependencies for a component that comes from a project dependency in a product's Bill of Materials.

— Described only in prose with no endpoint, UI path, or example

November 12, 2026 - Archived project mo…

└──▷ **BREAKING ON UPGRADE**

! `GET /api/public/v0/cves/{cveId}/metadata` and `GET /api/public/v0/findings/{projectVersionId}/{findingId}/cves` now return every `weaknesses[].description[].value` as a `CWE-<number>` string instead of a bare number

(e.g. `399` becomes `'CWE-399'`); the placeholders `CWE-0` and `CWE--1` are omitted, and a CVE whose only weakness was a placeholder now returns an empty `weaknesses` list.

WAS THIS USEFUL? 👍 👎

fickling

Feedback ⓘ CVE-2026-22612 CVE ↗ STABLE

Fickling is a Python library that manipulates, analyzes, and deserializes pickle files, helping detect and create malicious pickle payloads for security research.

`CVE-2026-22612`: Detection bypass due to `builtins` blindness (CVSS 7.8, affects `<= v0.1.6`)

▾ **3 more releases in this issue** · 2025-12-15 → 2026-01-09

> CVE-2026-22606 CVE ↗ STABLE
>
> `CVE-2026-22606`: Bypass via runpy.run_path() and runpy.run_module() (CVSS 7.8, affects `<= v0.1.6`)

> CVE-2025-67747 CVE ↗ STABLE
>
> `CVE-2025-67747`: Bypass via marshal.loads() and types.FunctionType() (affects `<= 0.1.5`)

> CVE-2025-67748 CVE ↗ STABLE
>
> `CVE-2025-67748`: Bypass via pty.spawn() (affects `<= 0.1.5`)

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

ecapture

Feedback ⓘ v2.11.0 NOTES ↗ RANK

eCapture is an eBPF-based tool that captures plaintext network traffic, SSL/TLS data, and bash commands without requiring CA certificates or code modifications.

> eCapture v2.11.0 adds richer process context to captured traffic and switches TLS modules to pcapng by default for out-of-the-box HTTP reconstruction.

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

01 **Pcapng default for TLS text-mode output** BREAKING 70

The default text-mode output source for the `tls`, `gotls`, and `gnutls` modules is now pcapng instead of the previous default, so reconstructed HTTP messages are the out-of-the-box experience; existing scripts or workflows that relied on the old default text source may behave differently.

— Clear before/after and named modules, but no migration command given v2.11.0

02 **Process identity captured and exposed in pcapng/ecaptureq**

NEW 65

eCapture now captures up to 256 bytes of process arguments per packet in pcapng output, giving analysts fuller context on which process generated each record. It also preserves the PID and Comm (process name) fields through the pcapng decryption pipeline so decrypted traffic retains its originating process identity, and the `ecaptureq` query interface now exposes these process arguments alongside TLS handshake metadata.

— Names fields and interface but no exact commands or query syntax v2.11.0

└──▷ **BREAKING ON UPGRADE**

! The default text-mode source for the `tls`, `gotls`, and `gnutls` modules is now pcapng instead of the previous default; existing scripts or workflows that relied on the old default text source may behave differently.

WAS THIS USEFUL? 👍 👎

Ettercap

Feedback ⓘ CVE-2026-125406 CVE ↗ STABLE

Ettercap is an open-source network security tool for man-in-the-middle attacks, enabling live packet sniffing, filtering, and protocol analysis on local networks.

`CVE-2026-125406`: Stack buffer underflow in the telnet dissector's backspace handling (CVSS 8.8, affects `<= 0.8.4.1`)

▾ **6 more releases in this issue** · 2026-11-06

> CVE-2026-125405 CVE ↗ STABLE
>
> `CVE-2026-125405`: Heap out-of-bounds read in the HTTP dissector's NTLM Type 3 parser (CVSS 8.1, affects `<= 0.8.4.1`)

> CVE-2026-125404 CVE ↗ STABLE
>
> `CVE-2026-125404`: size_t underflow in ccmp_decrypt leads to unbounded stack write (CVSS 7.5, affects `<= 0.8.4.1`)

> CVE-2026-125402 CVE ↗ STABLE
>
> `CVE-2026-125402`: Heap out-of-bounds read parsing the TLS SNI extension during SSL MITM (CVSS 8.1, affects `<= 0.8.4.1`)

> CVE-2026-125403 CVE ↗ STABLE
>
> `CVE-2026-125403`: Out-of-bounds read and unchecked memmem result in sslw_remove_sts (CVSS 8.1, affects `<= 0.8.4.1`)

> CVE-2026-125401 CVE ↗ STABLE
>
> `CVE-2026-125401`: Heap buffer overflow in the Gadu-Gadu dissector (gg_get_version) (CVSS 8.1, affects `<= 0.8.4.1`)

> CVE-2026-125400 CVE ↗ STABLE
>
> `CVE-2026-125400`: Use-after-free in PCRE2 filter substitution (func_pcre) (CVSS 8.8, affects `<= 0.8.4.1`)

WAS THIS USEFUL? 👍 👎

◆ Email Security

Rspamd

Feedback ⓘ 4.3.0 NOTES ↗ RANK

Rspamd is an open-source spam filtering system that analyzes emails using rules, statistical methods, and machine learning to assign spam scores.

> Rspamd 4.3.0 ships new detection rules for DKIM-signer abuse, tracking pixels, and far-future dates, adds SMTPUTF8 address syntax checks, and expands archive and truncated-content inspection — alongside a long list of breaking changes to SPF evaluation order, DKIM validation, symbol option encoding, and several Lua API functions.

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

01 **Archive inspection improvements** NEW 81

`archive.unpack` gains new `ignore_case` and `filters = false` options to control unpacking behavior. `EXE_IN_ARCHIVE` now also checks executable names inside archives named as documents, inside OOXML/ODF packages, and across the first 1000 archive entries.

Unpack an archive while ignoring case in filenames and disabling built-in libarchive filters, for controlled inspection of attachment contents.

```lua
archive.unpack(task, part, {ignore_case = true, filters = false})
```

— Runnable Lua call with exact option names given 4.3.0

02 **SMTPUTF8 address syntax checks** BREAKING 78

Rspamd now validates SMTPUTF8 addresses per draft-ietf-mailmaint-smtputf8-syntax, scoring ASCII-syntax violations as `INVALID_FROM_ASCII_ADDR` and `INVALID_RCPT_ASCII_ADDR` (weight 1.0). The existing `INVALID_FROM_8BIT` and `INVALID_RCPT_8BIT` symbols now fire only for genuinely 8-bit addresses, so existing rules or score overrides targeting the old symbols need updating.

— Names exact symbols and spec but no config example 4.3.0

03 **Truncated and incomplete content detection** NEW 75

New symbols `MIME_UNREADABLE_ARCHIVE`, `MIME_TRUNCATED_ARCHIVE`, `TEXT_WORDS_TRUNCATED`, `OOXML_INCOMPLETE`, `SVG_INCOMPLETE`, `PDF_NO_OBJECTS`, `PDF_DECODE_LIMIT`, `PDF_UNDECODED_STREAMS`, `PDF_PARSE_ERRORS`, and `PDF_TEXT_TRUNCATED` report content that was cut short,

incomplete, or unreadable across MIME, OOXML, SVG, and PDF documents. Text part `get_stats` also gains `words_sampled` and `words_truncated` fields to back this detection.

— All ten symbols and two API fields named verbatim `4.3.0`

04 **Symbol option encoding change** BREAKING 75

Symbol options are now stored verbatim; the characters `[]{};,` and `%` are percent-encoded only when options are joined, affecting `X-Spamd-Result`, `X-Virus`, `metadata_exporter`, ClickHouse output, and the symbols log line. This changes the encoding of previously written option values.

— All affected surfaces named, no migration tool given `4.3.0`

05 **Lua API behavior changes** BREAKING 73

`lua_util inflate` now returns errors as a second value instead of a single return. `lua_url get_redirected` is no longer an alias of `get_phished`, so code relying on it will silently change behavior. The image flag on URLs is now set only when a URL is seen solely as an image source, changing what `get_urls` returns when iterating. A `get_date` result of `0` now means `no date` in the time('message') selector and the spamassassin plugin, so `DATE_IN_PAST` rules no longer match mail without a `Date` header.

— Four named Lua functions and exact behavior changes `4.3.0`

06 **SPF term evaluation order change** BREAKING 70

SPF records are now evaluated in the order of their terms, and the qualifier of an `include` or `redirect` is applied after the mechanisms. Records like `-ip4:1.2.3.4 +ip4:1.2.3.0/24 -all` that previously allowed an address to pass will now reject it.

— Concrete before/after example record given `4.3.0`

07 **DKIM signature validation hardening** BREAKING 65

DKIM signatures with duplicated tags or `l=0` over a body are now rejected, as are RSA keys under 1024 bits and keys whose type does not match the `k=` tag. The `d=` domain is now lowercased consistently in results, symbol options, and glob maps.

— Clear before/after but no migration command `4.3.0`

08 **BAD_DKIM_SIGNER rule** NEW 60

New multimap rule `BAD_DKIM_SIGNER` matches every checked DKIM `d=` signer domain against the `bad_dkim_signers` glob map.

— Names rule and map but no setup steps `4.3.0`

09 **Archive output and file-counting limits** BREAKING 60

get_files(n) and get_files_full(n) now count files rather than names, with other names of an entry flagged as aliases. `archive.unpack` with no explicit limit is now capped at 1 GiB of output and one million headers.

— Named functions and limits but no migration snippet 4.3.0

10 **Statistics dump filtering by user** NEW 60

The `statistics_dump` command gains a new `dump --user` option to filter statistics dumps by user.

— Exact flag named, ready to run 4.3.0

THINNER COVERAGE BELOW

11 **HTML tracking pixel detection** NEW 50

New rule `HTML_TINY_PIXEL_TRACKER` and composite rule `HTML_TINY_PIXEL_PROBE` detect contentless messages carrying tiny external tracking pixels.

— Names rules, no scoring or map detail given 4.3.0

12 **Phishing URL mismatch rewrite trust maps** BREAKING 45

Phishing detection adds explicit rewrite trust maps for URL mismatch scoring. URL mismatches are now scored separately from phishing, with rewrite permissions sourced from dynamic maps, so existing phishing score tuning may need adjustment.

— Mechanism described but maps not named 4.3.0

13 **Far-future Date header detection** NEW 35

New rule `DATE_FAR_FUTURE_NO_TZ` flags `Date` headers set far in the future that lack a timezone.

— Single named rule with no scoring detail 4.3.0

└──▷ **BREAKING ON UPGRADE**

! URL mismatches are now scored separately from phishing; rewrite permissions come from dynamic maps, so existing phishing score tuning may need adjustment.

! Symbol options are now stored verbatim; `[]{};,` and `%` are percent-encoded only when options are joined (X-Spamd-Result, X-Virus, metadata_exporter, ClickHouse, the symbols log line), which changes the encoding of previously written option values.

! SPF records are now evaluated in the order of their terms, and the qualifier of an `include` or `redirect` is applied after the mechanisms — records like `-ip4:1.2.3.4 +ip4:1.2.3.0/24 -all` that previously allowed an address to pass will now reject it.

! DKIM signatures with duplicated tags or `l=0` over a body are now rejected, as are RSA keys under 1024 bits and keys whose type does not match the `k=` tag. The `d=` domain is lowercased in results, symbol options, and glob maps.

! SMTPUTF8 ASCII syntax violations now score as `INVALID_FROM_ASCII_ADDR` and `INVALID_RCPT_ASCII_ADDR` (1.0); `INVALID_FROM_8BIT` and `INVALID_RCPT_8BIT` fire only for 8-bit addresses — existing rules or score overrides targeting the old symbols will need updating.

! get_files(n) and get_files_full(n) now count files rather than names; other names of an entry are flagged as aliases. `archive.unpack` with no limit is now capped at 1 GiB of output and one million headers.

! `lua_util inflate` now returns errors as a second value instead of a single return.

! `lua_url get_redirected` is no longer an alias of `get_phished` — code relying on this alias will silently change behavior.

! The image flag on URLs is now set only when a URL is seen solely as an image source; `get_urls` now returns a link that is also an image source, changing iteration results.

! A `get_date` result of `0` now means `no date` in the time('message') selector and the spamassassin plugin, so `DATE_IN_PAST` rules no longer match mail without a `Date` header.

! PDF content stream parsing, page ordering, and word spacing changed; extracted text from PDFs will differ from previous versions.

WAS THIS USEFUL? 👍 👎

◆ IDS / IPS

maltrail

Feedback ⓘ 3.6 NOTES ↗ RANK

Maltrail is a malicious traffic detection system that uses blacklists and heuristics to identify suspicious network activity and potential threats.

> Maltrail 3.6 adds ERSPAN and MPLS packet decapsulation, new DNS CNAME, ICMP-tunneling and IDN-homograph detection heuristics, a trail-store memory-map cache that cuts cold reloads from 0.78s to 0.08s, and a new sensor health and alert-persistence layer, alongside breaking config key renames in `maltrail.conf`.

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

01 Trail-store memory-map cache NEW

95

A new memory-map cache for the trail table can be enabled or disabled with `TRAIL_STORE_CACHE` in `maltrail.conf`, and its integrity can be verified with `maltrail-sensor --verify-trail-store`. On 1.9M trails it reduces a cold `-T` reload from 0.78s to 0.08s and peak RSS from 145MB to 89MB.

Verify the trail-store cache is consistent with the current trails.csv before relying on the fast mapped-load path.

```
$ maltrail-sensor --verify-trail-store
```

— Named config key, CLI flag and measured before/after numbers 3.6

02 Sensor health monitoring and dashboard reporting NEW

90

Adds sensor health and alert-persistence infrastructure: a configurable stale threshold via `SENSOR_STALE_AFTER` in `maltrail.conf` (default 15 minutes) with stale/recovery alerts logged to `error.log`; a configurable remote clock-skew tolerance via `LOG_SERVER_SKEW` in `maltrail.conf` that counts and shows out-of-window senders instead of silently misfiling them; a heartbeat file at `LOG_DIR/<sensor>.heartbeat` for single-host installs without `LOG_SERVER`; persistent sensor state in `LOG_DIR/sensors.json` so a server restart retains knowledge of which sensors should be reporting; a persistent alert queue in `LOG_DIR/alerts.json` preserving up to 1,000 pending alerts and the tail position across restarts; and the `/sensors` endpoint now also shows webhook delivery status and, sourced from packet counts carried in the heartbeat, dropped-packet share per sensor, with drop alerting tunable per sensor via `SENSOR_DROP_ALERT` and `SENSOR_DROP_WINDOW` in `maltrail.conf`.

Tune drop alerting on a busy sensor so the server notifies once per hour when more than 1% of packets are dropped.

```ini
SENSOR_DROP_ALERT = true
SENSOR_DROP_WINDOW = 3600
```

— Five named config keys and file paths, no single runnable command Names config keys with a concrete tuning example

3.6

03 **Typed TLS fingerprint trails and data validation tooling** NEW

85

Static trail files now support typed TLS fingerprints per line (`ja3:<md5>` , `ja4:<ja4>` , `cert:<sha1 or sha256>`), with `check_trails.py` validating each type and certificate trails accepting SHA-256 as well as SHA-1. A new format gate `sensor/tools/check_data.py` validates every `data/*.txt` file, catching loader-silent typos, unsupported inline comments, and IPv6 ranges, to be run on any `data/` change.

— Names formats and both validator scripts to run 3.6

04 **Config key renames and capture defaults in maltrail.conf**

85

BREAKING

`CAPTURE_FANOUT` is renamed to `CAPTURE_WORKERS` and `CHECK_TLS_CERTIFICATES` is renamed to `CHECK_TLS_FINGERPRINTS` in `maltrail.conf` ; old names are still read but should be updated, and `PROCESS_COUNT` is no longer set. `CAPTURE_BUFFER` is no longer mandatory — the default 10% libpcap ring is used if unset, though existing explicit values are kept. The default `CAPTURE_FILTER` now restricts QUIC to long-header (handshake) packets only, while existing custom filters are kept until manually updated.

— Names every renamed/changed config key with migration behaviour 3.6

05 **Configurable alert throttle** IMPROVED

73

Alert throttle is now configurable via `ALERT_THROTTLE` in `maltrail.conf` (default `86400` , one message per threat per day); the `%(threat)s` template variable now carries the dashboard threat ID.

— Named config key and default, no example command given 3.6

06 **Configurable dashboard session expiration** NEW

70

A new `SESSION_EXPIRATION_HOURS` config key in `maltrail.conf` controls how long a dashboard session stays valid, defaulting to 24 hours.

Reduce dashboard session length to 8 hours for analyst accounts on shared workstations.

```ini
SESSION_EXPIRATION_HOURS = 8
```

— Named config key with default and a runnable example 3.6

07 **Dashboard day-loading via Web Worker** IMPROVED 70

Dashboard day-loading now runs in a Web Worker (`js/worker.js`) against the same engine as the main page (`js/engine.js`), keeping the table interactive during load; the in-thread path can be forced with `localStorage.mt_worker = "0"` .

— Names files and the override flag to disable it 3.6

08 **ERSPAN and MPLS packet decapsulation** NEW 68

The sensor now decapsulates ERSPAN type I and type III (including the optional platform sub-header), MPLS label stacks up to 8 deep, and Ethernet over MPLS with or without a control word.

— Mechanism detailed but no config surface to act on 3.6

09 **New detection heuristics** NEW 68

DNS responses are now checked against trail tables for listed CNAME targets, naming both the queried and aliased domain when a client resolves an innocent name aliased to a listed one; a new `icmp_tunneling` heuristic detects ICMP-based tunnels (icmpsh, ptunnel, exfiltration) by flagging more than 10 distinct echo payloads from one host to one destination that match no known ping-tool pattern; and a new `idn_homograph` heuristic detects IDN homograph attacks where an `xn--` label with Cyrillic, Greek, or Armenian lookalike letters resolves to a whitelisted domain it is not.

— Clear behaviour described, no tunable surface named Names heuristic and its exact trigger threshold Names heuristic and the scripts it checks

3.6

THINNER COVERAGE BELOW

10 **Reverse DNS lookup action in dashboard** NEW 45

The dashboard adds a right-click Reverse DNS action on event addresses, resolved by the server on demand.

— Clear UI path but no further mechanism given 3.6

└──▷ **BREAKING ON UPGRADE**

! `CAPTURE_FANOUT` is renamed to `CAPTURE_WORKERS` and `CHECK_TLS_CERTIFICATES` is renamed to `CHECK_TLS_FINGERPRINTS` in `maltrail.conf` ; old names are still read but

should be updated. `PROCESS_COUNT` is no longer set.

! `CAPTURE_BUFFER` is no longer mandatory in `maltrail.conf`; the default 10% libpcap ring is used if unset, so existing configs that set it explicitly keep their value unchanged but new installs omit it.

! An existing `maltrail.conf` with a custom `CAPTURE_FILTER` keeps its current filter until manually updated; the new default filter restricts QUIC to long-header (handshake) packets only.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ SaaS Security Posture

Nudge Security

Feedback ⓘ changelog · Nov 14, 2026

CHANGELOG ↗

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

Nudge Security adds breached password detection with in-browser remediation nudges and AI agent verdict filtering for browser extensions.

NEW DETECTION Configure breached password detection via browser extension settings to detect when employees log in with passwords exposed in known data breaches, trigger in-browser nudges to prompt password changes, and surface affected apps and accounts on a new password risk findings card in the Identity & location dashboard.

Filter the browser extension table by AI agent verdict — malicious, compromised, risky, or safe — to compare verdicts against each extension's approval status, such as finding malicious extensions that are still approved.

UPDATED DETECTION Region-specific sensitive data detection in AI conversations

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ changelog · Nov 14, 2026 CHANGELOG ↗

RANK

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

> Sysdig Secure's Registry Scanner (v0.14.0) adds credential-free authentication to Azure Container Registry via managed or workload identity and extends registry scanning to Red Hat OpenShift Service on AWS (ROSA) with hosted control planes reading Amazon ECR.

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

01 **ACR managed/workload identity authentication in Registry Scanner**

NEW 77

Registry Scanner v0.14.0 authenticates with Azure Container Registry (ACR) using managed identity or workload identity when running on AKS — via the cluster's kubelet identity or a user-assigned identity federated to its service account — removing the need to store registry credentials. Both RBAC and RBAC + ABAC repository permission modes are supported, and new log messages distinguish ACR identity authentication failures between an identity that cannot get a token, one with no role on the registry, and a request blocked by registry network rules.

— Names mechanism, permission modes and failure cases but no config keys or flags.

changelog-20261114-348548da

02 **ROSA hosted control plane support in Registry Scanner** NEW

62

Registry Scanner v0.14.0 now supports Red Hat OpenShift Service on AWS (ROSA) with hosted control planes, reading Amazon Elastic Container Registry (ECR) via an IAM role bound to its service account, so AWS access keys no longer need to be stored in the cluster.

— Clear platform and mechanism but no setup steps or command given.

changelog-20261114-348548da

WAS THIS USEFUL? 👍 👎

Turbot Steampipe

Feedback ⓘ v2.5.0 NOTES ↗ CODE ↗ RANK

Steampipe is an open-source tool that lets you query cloud infrastructure, APIs, and security data using SQL.

> Steampipe v2.5.0 upgrades the embedded PostgreSQL to 14.24.0 and tightens its authentication model, relocating the unix socket and locking down root access and `pg_hba.conf` handling.

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

01 **pg_hba.conf management hardened** BREAKING 75

Custom host-based authentication rules can be placed in `pg_hba.conf.d` to be applied first on every start. Steampipe now writes `pg_hba.conf` before every startup and never overwrites a file it did not create, stopping with an error if it encounters a `pg_hba.conf` it did not write, requiring manual resolution before startup proceeds.

— Names exact files and gives a concrete fallback path for custom rules v2.5.0

02 **Root database access locked to unix socket** BREAKING 65

Steampipe now restricts `root` database connections to the OS user who launched it, via the unix socket only; any TCP or other-user `root` connections that previously worked will be rejected.

— Clear mechanism and scope, no example or migration command given v2.5.0

03 **Unix socket relocated to data directory** BREAKING 65

The PostgreSQL unix socket now lives in the data directory instead of `/tmp`, with a clear error and guidance if the resulting path exceeds the OS limit; external clients or scripts that hardcode the `/tmp` socket path will fail to connect.

— Names old and new socket locations and the failure condition v2.5.0

THINNER COVERAGE BELOW

04 **PostgreSQL upgraded to 14.24.0 with data preservation** IMPROVED 55

Steampipe now upgrades the embedded PostgreSQL to `14.24.0`, preserving the existing data folder on same-major upgrades and carrying over the newest older database.

— Describes upgrade mechanism but no user action required v2.5.0

05 **Plugin SDK and FDW dependency updates** IMPROVED 30

Steampipe now updates `steampipe-plugin-sdk` to v6 and the embedded `steampipe-postgres-fdw` to `v2.3.0`.

— Bare version bump with no explanation of impact `v2.5.0`

└──▷ **BREAKING ON UPGRADE**

! The PostgreSQL unix socket path has moved from `/tmp` to the data directory; external clients or scripts that hardcode the `/tmp` socket path will fail to connect.

! Only the OS user who runs Steampipe can now connect as `root`, and only through the unix socket; any TCP or other-user `root` connections that previously worked will be rejected.

! Steampipe now writes `pg_hba.conf` before every start and stops with an error if it encounters a `pg_hba.conf` it did not write; a manually managed `pg_hba.conf` will block startup until resolved.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Entitlements is now Identity security: NOTES ↗

RANK

Plerion is a cloud security platform that provides comprehensive risk management, vulnerability detection, and threat prioritization across cloud environments and workloads.

> Plerion consolidates identity-related capabilities into a new Identity security section, adding a cross-cloud Overview and reorganizing access grant review into a single Access grants tab that replaces the old per-cloud overview tabs.

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

01 Identity security Overview section NEW

75

Plerion adds a new Identity security section at `/identity-security` that surfaces cross-cloud identity posture across AWS, Azure, and GCP. The Overview tab shows account and identity counts, open findings with 30-day trends, a 7-day change summary, risk by account, and a prioritized findings list, with a cloud selector to filter by provider.

Review your cross-cloud identity posture and prioritized open findings in one place.

> 📍 In the side navigation, go to Identity security › Overview. Use the cloud selector to filter to AWS, Azure, or GCP individually, then check Identity risk by account and the prioritized findings list.

— Names the exact route and metrics shown, with navigation steps.

Entitlements is now Identity security:

02 Access grants tab with external exposure views BREAKING

75

The Access grants tab, now at `/identity-security/access-grants`, lets users review resource access grants, with new External and Untrusted external views for scoping external access exposure. This retires the separate AWS, Azure, and GCP overview tabs, including the identity risk trend, identity top risks, and the AWS External access card, consolidating external access review into this single tab.

Investigate external access exposure across your cloud environments after the AWS External access card retirement.

> 📍 In the side navigation, go to Identity security › Access grants, then select the External or Untrusted external view to scope and review externally accessible resource grants.

— Names the route, the retired tabs, and replacement views with steps.

Entitlements is now Identity security:

└──▷ **BREAKING ON UPGRADE**

! The separate AWS, Azure, and GCP overview tabs (with identity risk trend, identity top risks, and the AWS External access card) have been retired; external access is now reviewed via the External and Untrusted external views on the Access grants tab at `/identity-security/access-grants`.

WAS THIS USEFUL? 👍 👎

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Nov 10, 2026 BLOG ↗ RANK

Rapid7 InsightCloudSec is a cloud-native security platform that provides continuous compliance, threat detection, and automated remediation across multi-cloud environments.

> Rapid7 InsightCloudSec's November release expands cloud inventory coverage and adds new posture insights and search filters.

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

01 **New cloud resource coverage, Insights, and Query Filters** NEW

30

InsightCloudSec now supports `Amazon SageMaker Ground Truth` resources for cloud inventory and analysis, and the release adds new Insights for cloud security posture checks along with new Query Filters for narrowing cloud resource searches.

— Names one specific resource type but no concrete Insight or filter names given.

2026-11-10-5f389534

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

kubescape

Feedback ⓘ v4.0.21 NOTES ↗ RANK

Kubescape is an open-source Kubernetes security platform that scans clusters, YAML files, and Helm charts for misconfigurations and vulnerabilities against security frameworks.

> Kubescape's v4.0.21 release focuses on scan customization and observability, adding ValidatingAdmissionPolicy-based custom rules, a file watcher for continuous scanning, and refined evidence output.

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

01 **ValidatingAdmissionPolicy support in custom rules** NEW 60

`kubescape scan` now accepts user-authored ValidatingAdmissionPolicy YAML via the `--custom-rules` flag.

— Names flag and command, clear usage path. v4.0.21

THINNER COVERAGE BELOW

02 **Test case filtering and JUnit XML reports** NEW 55

Adds `--run` to filter which policy test cases run, plus support for writing policy test results as a JUnit XML report.

— Names the flag and output format but little mechanism detail. v4.0.21

03 **Evidence output improvements for cluster scans** IMPROVED 55

Prints a `kubectl` command for each evidence path on cluster scans, and `--show-evidence` now shows current values for delete paths.

— Names the flag and output change but limited mechanism. v4.0.21

04 **File watcher for kubescape watch** NEW 50

Adds a file watcher with a debounced, sequential scan scheduler for `kubescape watch`.

— Describes scheduling mechanism and names the command. v4.0.21

05 **verify_remediation tool in MCP server** NEW 40

Adds a `verify_remediation` tool to the MCP server for local IaC remediation verification.

— Names the tool but no mechanism or usage detail. v4.0.21

06 **Kustomize resource attribution to source file** IMPROVED

35

Attributes each Kustomize-rendered resource to the source file Kustomize read it from.

— States the behavior change but no further specifics. v4.0.21

07 **JSON schema for fleet report** NEW

25

Adds a JSON schema for the fleet report.

— Bare mention with no further detail. v4.0.21

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Trend Vision One Feedback ⓘ

AI Secure Access supports content inspection for the Gemini add-on (2026-11-15)

BLOG ↗

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

NEW INTEGRATION Gemini add-on content inspection for Google Workspace

WAS THIS USEFUL? 👍 👎

◆ AI SOC

Exaforce

Feedback ⓘ

Force ER: Emergency incident response in 15 minutes, with no retainer

NOTES ↗

RANK

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

> Exaforce launched Force ER, an on-demand emergency incident response service with a 15-minute responder connection and 24-hour environment onboarding, backed by automated log replay, a shared War Room, mobile incident controls, and an optional retainer tier.

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

01 **Force ER on-demand emergency incident response** NEW

80

Exaforce adds Force ER, an emergency incident response service available to any organization under active attack with no existing customer relationship or retainer required. It connects a live responder within 15 minutes and goes live in the customer's environment within 24 hours; the incident hotline is `+1 (844) 660-1663` or exaforce.com/incident-response.

— Names concrete hotline and URL, exact time thresholds

Force ER: Emergency incident response i…

THINNER COVERAGE BELOW

02 **Automated 30-day log replay for incident timelines** NEW

55

As soon as an organization connects to Force ER, Exaforce replays the last 30 days of identity, cloud, SaaS, code, and VPN logs to build a pre-built incident timeline without manual log correlation.

— Clear mechanism and scope but no direct user action

Force ER: Emergency incident response i…

03 **Mobile incident response via ExaGo** IMPROVED 55

Exaforce's ExaGo mobile interface now lets responders approve containment actions, ask incident-status questions, and join the War Room from a phone or Apple Watch during an active incident.

— Names device surfaces and actions but no exact app path

Force ER: Emergency incident response i…

04 **Shared War Room for incident collaboration** NEW 45

Force ER adds a shared War Room where the customer's team and the MDR team work from the same live incident view, with communication synced to Slack and findings routed to named stakeholders.

— Describes collaboration surface but no setup detail given

Force ER: Emergency incident response i…

05 **Optional incident response retainer tier** NEW 45

Exaforce adds an optional incident response retainer tier that pre-connects data sources, establishes named contacts, and provides a contractual SLA so investigation begins on the first call.

— Describes tier benefits but no pricing or onboarding specifics

Force ER: Emergency incident response i…

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Test v3.0.0: Testing with Tenzir 7 (2026-11-14) BLOG ↗ RANK

Tenzir is a data pipeline engine for security teams that collects, transforms, enriches, and routes telemetry.

> Tenzir's headline release this window is v7.0.0-rc.1, adding an MCP server and built-in assistant for AI agents, versioned packages with live rollout, and breaking changes to several source operators, while companion tools pushed pushdown-based query acceleration for SQL/Splunk/SentinelOne sources, Iceberg sink efficiency, Tenzir Test's move to Tenzir 7, and TQL editor tooling updates.

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

01 **Tenzir Test runs on Tenzir 7 via tenzir run** NEW 80

Tenzir Test v3 runs test pipelines using the `tenzir run` command from Tenzir 7, fetches a package's locked dependencies before running its tests, and automatically injects project fixtures into Python tests without requiring explicit imports.

— Names exact command and mechanism, no config example

Tenzir Test v3.0.0: Testing with Tenzir…

02 **Query pushdown into DuckDB, ClickHouse, MySQL, MSSQL, and Splunk**

UPDATED SUPPORTED PLATFORM UPDATED COMPLIANCE FRAMEWORK NEW INTEGRATION UPDATED INTEGRATION NEW COMPLIANCE FRAMEWORK IMPROVED 80

Tenzir now pushes filters, sorting, and limits after summaries down into SQL databases (DuckDB, ClickHouse, MySQL, Microsoft SQL Server) and Splunk, reducing the data returned over the wire. Splunk searches now return typed results rather than raw strings for faster downstream processing, and pushdown now reaches through more stateful operators, `ocsf_cast`, and field selection functions in pipelines.

— Names five backends and specific operators affected by pushdown

Tenzir v6.21.0: Faster DuckDB, ClickHou…

03 **Tenzir Test v3 breaking changes for Tenzir 6 projects** 75

BREAKING

The built-in `node` fixture is removed; projects still on Tenzir 6 should pin `tenzir-test` below version 3, since Tenzir Test v3 requires Tenzir 7 and will not work with

Tenzir 6.

— Gives exact migration step and version pin Tenzir Test v3.0.0: Testing with Tenzir…

THINNER COVERAGE BELOW

05 **Automatic package.yaml version bump on release** IMPROVED

50

Tenzir Ship releases now automatically update the version field in a Tenzir package's `package.yaml`, removing the need for manual manifest edits after cutting a release.

— Names config file but change is automatic, no user action

Tenzir Ship v2.4.0: Automatic version u…

06 **Tree-sitter TQL grammar adds editor highlighting and structure queries**

IMPROVED

50

Tenzir's Tree-sitter TQL grammar adds bracket pair and verbatim indentation queries for editor integrations, and distinct highlight captures so editor themes can color pipeline operators, times, durations, IP addresses, and subnets apart from function calls and other literals.

— Names grammar features but no setup or config example

Tree-sitter TQL v1.4.0: Distinct highli… Tree-sitter TQL v1.3.0 (2026-11-13)

08 **Fewer small files from to_iceberg writes** IMPROVED

45

Tenzir's `to_iceberg` operator now writes fewer small data files, reducing storage fragmentation for Iceberg table sinks.

— Names the operator but no size thresholds or config given

Tenzir v6.20.1: Leaked CA keys no longe…

09 **Deeper filter pushdown to SentinelOne** NEW INTEGRATION IMPROVED

45

Tenzir now pushes more filter conditions inside disjunctions down to SentinelOne, reducing data transferred from the SentinelOne API.

— Names SentinelOne API but no example query or limit

Tenzir v6.20.1: Leaked CA keys no longe…

NEW INTEGRATION Redis source and sink operators

NEW DETECTION Suricata schema types for NTP, LDAP, LLMNR, DNP3, drop, and frame events

NEW INTEGRATION Write metrics to Google Cloud Monitoring

NEW INTEGRATION Cloudflare Workers AI and OpenAI Decisions endpoints for

ai_decide

NEW INTEGRATION Windows Event Forwarding without a Windows collector

└──▷ BREAKING ON UPGRADE

! The built-in `node` fixture is removed; projects on Tenzir 6 should pin `tenzir-test` below version 3.

! Tenzir Test v3 requires Tenzir 7; it will not work with Tenzir 6.

! Native Hydrolix semantics are now applied to operators after `from_hydrolix`, changing prior behavior.

! Native Microsoft Graph semantics are now applied to operators after `from_microsoft_graph`, changing prior behavior.

! Native CloudWatch Logs semantics are now applied to operators after `from_amazon_cloudwatch`, changing prior behavior.

! Content-skipping filters and limits are now applied after `from_microsoft_365_activity`, changing prior behavior.

! Operators after `from_velociraptor` now stay in the pipeline by design, changing prior exit behavior.

! Sink byte metrics now count the delivered payload, not the raw input, changing prior metric values.

`Tenzir v6.22.0: OCSF compaction and native source semantics (2026-11-12)`

BLOG ↗ STABLE

Tenzir adds `ocsf_compact` for lossless OCSF compaction, `category_uid` / `type_uid` computation in `ocsf_derive`, a `version` function, time-zone support in `parse_time`, automatic syslog framing detection, and lower memory and CPU usage.

Add lossless OCSF compaction with the `ocsf_compact` operator to reduce event size without losing fields.

Compute `category_uid` and `type_uid` automatically in the `ocsf_derive` operator.

Support time zones in the `parse_time` function.

Add a `version` function to retrieve the current Tenzir version inside a pipeline.

Automatic syslog framing detection added to the syslog parser.

+1 more

└──▷ BREAKING ON UPGRADE

! Several OCSF operators are deprecated in this release — existing pipelines using the deprecated forms may need updating.

Tenzir v6.20.0: A new data model for faster pipelines (2026-10-26) BLOG ↗

STABLE

Tenzir adds a new heterogeneous data model for faster pipelines, typed AI decisions via `ai_decide`, secrets in function arguments, Parquet encryption, and new built-in functions including `type_kind`, Luhn, and CRC-32.

Enable typed decisions in pipelines with the new `ai_decide` operator, which now supports Decisions endpoints and typed output.

Pass secrets as function arguments and use secret keys with the `hmac` function, with secrets automatically censored on assignment.

Compute Luhn check digits and CRC-32 checksums with new built-in functions.

Encrypt Parquet output with a shared AES key.

Inspect the kind of a type at runtime with the new `type_kind` function.

+8 more

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ More granular control over product access (2026-11-13)

BLOG ↗ RANK

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

> Tines added per-product-area access controls, letting teams on the same workspace hold different permissions across product areas.

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

01 **Per-product-area access permissions** NEW 33

Tines now lets admins control access to each product area independently, so analysts and builders who share a team can hold different product permissions rather than a single uniform access level.

·

– Describes capability but no exact UI path or setting names

More granular control over product acce…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Importing Power BI and Tableau files with Genie Code will soon be available by default f

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 previewed Power BI/Tableau import via Genie Code for compliance workspaces and introduced fine-grained Unity Catalog permissions for pipeline flows.

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

01 **Flows as securable objects in Unity Catalog** NEW 85

Pipeline flows can now be individually secured in Unity Catalog (Beta) using fine-grained privileges: `READ FLOW`, `REFRESH FLOW`, `WRITE FLOW`, `APPLY TAG`, and `MANAGE` on flows, or `CREATE FLOW` on schemas.

Restrict a service principal to only triggering a specific pipeline flow refresh without granting broader pipeline access.

```
$ GRANT REFRESH FLOW ON FLOW my_catalog.my_schema.my_flow TO
`service-principal@example.com`;
```

— Names all five privileges plus a runnable GRANT command

Flows are securable objects in Unity Ca…

02 **BI file import via Genie Code** NEW 65

Power BI and Tableau files can be imported as AI/BI dashboards and metric views using Genie Code. It is available now for workspaces with the compliance security profile enabled by turning on Import BI files using Genie Code under Previews in workspace settings, and will be on by default starting December 15, 2026.

Enable Genie Code BI file import early for a compliance security profile workspace before the December 15, 2026 default rollout.

📍 In the workspace settings, go to Previews and enable Import BI files using Genie Code.

— Names exact setting and rollout date, but no mechanism detail

Importing Power BI and Tableau files wi…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

IT-Harvest

Feedback ⓘ announced · Nov 9, 2026 BLOG ↗ RANK

Data-driven analyst firm covering the entire cybersecurity industry, with 4,000+ vendors and 11,000+ products in one living research platform (the IT-Harvest Dashboard).

> IT-Harvest launched the CX-Index, a market-cap-weighted stock index tracking 23 cybersecurity vendors.

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

01 CX-Index stock index for cybersecurity vendors NEW

70

IT-Harvest added the CX-Index, a market-cap-weighted stock index tracking 23 cybersecurity vendors, modeled after the S&P 500 weighting methodology. It is viewable at `it-harvest.com/cx-index`.

— Names methodology and URL but no further mechanism detail launch-20261109-77ec4348

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

defectdojo

Feedback ⓘ 3.5.103 NOTES ↗

Defectdojo is an open-source Unified Vulnerability Management, DevSecOps & ASPM

DefectDojo 3.5.103 adds an Early Light parser; performance improvements speed up reimport JIRA lookups, tag indexing, and audit log history builds.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.5.103 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.5.103
```

NEW INTEGRATION DefectDojo adds an Early Light parser to import Early Light scan results via the standard importer pipeline.

Reimport now batches finding-group JIRA lookups in `process_groups_for_all_findings` to reduce round-trips.

Tag tables now index UPPER(name) so adding a tag no longer requires a full table scan.

Audit log history index builds now receive dedicated memory and workers to complete faster.

WAS THIS USEFUL? 👍 👎

gvm-libs

Feedback ⓘ v23.16.0 NOTES ↗ RANK

gvm-libs is a C library providing core functions and utilities for the Greenbone Vulnerability Management (GVM) framework and its components.

> gvm-libs v23.16.0 adds programmatic control over OSP connection timeouts and retries, and significantly expands Active Directory connector search capabilities.

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

01 **Active Directory connector search improvements** NEW

55

The Active Directory connector now supports searching container objects, extending the scope of AD queries beyond flat group and user lookups; it also supports group DN membership searches, and enforces base DN confinement checks to restrict the scope of AD searches.

— Names three distinct AD search behaviors but no API or config names. v23.16.0

02 **OSP connection timeout and retry controls** NEW 45

New getter and setter functions were added to the `osp` library to programmatically get and set the OSP connection timeout, and to get and set the maximum number of OSP connection retries.

— Describes the capability but no exact function names given. v23.16.0

WAS THIS USEFUL? 👍 👎

oxo

Feedback ⓘ v2.23.0 NOTES ↗ RANK

> oxo now captures developer fields in store listing fingerprints, extending the data collected during app-store asset scans.

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

01 **Developer fields in store listing fingerprints** IMPROVED

30

oxo now captures developer fields in store listing fingerprints, extending the data collected during app-store asset scans.

— Names the field addition but no schema, scope or usage detail. v2.23.0

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

N-able N-central

Feedback ⓘ

Tags management UX updates (2026-11-13) BLOG ↗ RANK

N-able N-central is a secure endpoints and automate IT operations at scale. N-central delivers unified endpoint management with advanced RMM capabilities for complex environments.

> N-able N-central shipped a preview of multi-level patch policy management with automatic inheritance across organisational tiers, alongside a reworked Tags management UX with scope filtering and a unified creation flow.

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

01 Multi-level patch policy management (preview) NEW

84

Patch policies now cascade automatically down the organisational hierarchy — Service Organisation, Customer, and Site — with more specific policies at lower levels overriding inherited ones when assets conflict. The policy summary page in policy details shows exactly which assets a policy is deployed to before you deploy an override, and a new Add tags button on the Assets targeting step lets you create tags on the fly during policy creation without leaving the workflow.

Override a Service Organisation baseline policy for a customer or site with unique patching requirements, without affecting the rest of your infrastructure.

> 📍 In the Policies grid, use the Scope Selector to switch from Service Organisation to the target Customer or Site, then create a new policy — it will automatically override the inherited SO-level policy only for assets at that scope.

Confirm which assets will be affected before activating a new policy override at the Customer or Site level.

> 📍 Open the policy, go to its policy summary page in policy details, and review the effective assets list to verify deployment scope before saving.

— Clear inheritance/override mechanism and UI steps, but still preview-stage.

Introducing Multi-level Policy Manageme…

02 Scope Picker and unified creation flow in Tags management

IMPROVED

68

The Tags grid now has a Scope Picker at the top that filters tags by organisation, showing tags inherited from parent levels and those available to child scopes. Tag creation and assignment are combined into one flow — choosing type, scope (from the org tree), and assets all within the creation form with no context switching — and the grid gains an `Assets` column showing total asset counts per tag, including assets outside the user's access scope.

When managing tags across a multi-tenant hierarchy, use the Scope Picker to narrow the Tags grid to your organisation before creating or assigning a tag.

> 📍 In N-central, go to Tags › Tags grid › Scope Picker (top of grid), select your organisation, then click `Create Tag` — set the type, scope, and asset assignment all within the single creation form.

— Names specific UI elements and workflow but no API/config surfaces given.

`Tags management UX updates (2026-11-13)`

WAS THIS USEFUL? 👍 👎

◆ Attack Surface Management

bbot

Feedback ⓘ announced · Nov 12, 2026 BLOG ↗ RANK

Bbot is the recursive internet scanner for hackers.

> BBOT 3.1 replaces massdns with a new BlastDNS resolution engine for roughly 40% faster subdomain scans, turns CDN/WAF port filtering on by default, and ships seven new modules alongside a refreshed nuclei integration — plus a string of breaking changes including a Python 3.11 requirement and a reworked preset system.

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

01 BlastDNS-powered DNS resolution and resolver control NEW

93

BBOT now uses BlastDNS for all DNS traffic including brute-forcing, replacing massdns; adaptive rate control cuts failure rates and finishes subdomain scans about 40% faster, limiting the resolver pool to 600 randomly chosen public resolvers per scan. The new `--custom-nameservers` flag lets you pin BBOT to specific resolvers for normal DNS resolution instead of the default public pool.

Pin BBOT to your internal resolvers for normal DNS resolution instead of using the default public pool — useful in lab or air-gapped environments.

```
$ bbot -t example.com --custom-nameservers 10.0.0.53 10.0.0.54 -p subdomain-enum
```

— Names mechanism, flag, and concrete 40%/600-resolver figures with runnable example.

launch-20261112-a0e1a802

02 Built-in CDN/WAF port filtering, replacing portfilter module

88

BREAKING

CDN/WAF port filtering is now built into BBOT and on by default, dropping non-80/443 ports on `cdn` - or `waf` `-tagged` hosts; revert with the new `--allow-cdn-ports` flag. The standalone `portfilter` module is removed, so it must be dropped from existing presets.

— Exact flag, module name, and migration step given. launch-20261112-a0e1a802

03 js_unpacker module for obfuscated JavaScript NEW

88

New `js_unpacker` module transparently unpacks source maps, webpack, obfuscator.io, Next.js, and Dean Edwards' packer JavaScript, feeding results back into BBOT for more endpoint and secrets discovery; it has been added to the `js-audit` preset.

Run a full JS audit against a target to unpack obfuscated bundles and surface leaked secrets with TruffleHog — `js_unpacker` is now included in the `js-audit` preset.

```
$ bbot -t example.com -p js-audit
```

— Names all unpacked formats, preset, and includes a runnable command.

launch-20261112-a0e1a802

04 **Nuclei integration: new presets, pass-through options, cached indexing**

IMPROVED 70

Adds nine new nuclei pass-through options and two new presets: `nuclei-recommended` (built on ProjectDiscovery's curated `recommended` profile) and `nuclei-new`. Nuclei template indexing is now cached, cutting startup time from about 35 seconds to near-instant.

— Names presets and startup-time improvement, but the nine options are unnamed.

launch-20261112-a0e1a802

05 **Preset system: module_dirs and preset condition changes**

70

BREAKING

`module_dirs` now only works as a top-level preset key, and BBOT prompts before loading each custom module folder (use `-y` to accept). Preset conditions are now read-only — they can warn or abort but can no longer modify the preset.

— Names config key and flag with clear before/after behavior. launch-20261112-a0e1a802

06 **New passive subdomain enumeration modules: otilabs and scanmalware**

NEW 60

Adds `otilabs` (queries OTI Labs, requires an API key) and `scanmalware` (draws from ScanMalware's public scan archive) as new passive subdomain enumeration modules.

— Module names and data sources given, but no enable/run example.

launch-20261112-a0e1a802

THINNER COVERAGE BELOW

07 **response_to_json() helper removed** DEPRECATED 55

helpers.response_to_json() is removed; use response_to_event_dict() instead.

– Clear migration swap naming both functions. launch-20261112-a0e1a802

---

08 **nowafpls module for WAF-bypass padding detection** NEW

50

New `nowafpls` module detects when the WAF-bypass padding technique is usable and is integrated into `lightfuzz` and other modules to test for vulnerabilities through WAFs.

– Names module and integration but gives no usage path. launch-20261112-a0e1a802

---

09 **Technology naming and event tagging changes** BREAKING

40

CPE-based technologies are now named by product, and in-scope events are no longer tagged `affiliate`.

– States the change but no example of old vs new output. launch-20261112-a0e1a802

---

10 **dnsdumpster now requires an API key** BREAKING

35

`dnsdumpster` now requires an API key to be used.

– Names module but no detail on how to supply the key. launch-20261112-a0e1a802

---

11 **waf_bypass scoped to targets by default** BREAKING

35

`waf_bypass` now only tests scan targets by default.

– Single-line scope change with no further mechanism. launch-20261112-a0e1a802

---

12 **BBOT User-Agent header on API requests** BREAKING

30

API requests now send a `BBOT/<version>` User-Agent header.

– Names the header format but no further context given. launch-20261112-a0e1a802

---

13 **excavate long-redirect detection** NEW

25

New `excavate` submodule detects long-redirects.

– Bare one-line mention with no mechanism or usage. launch-20261112-a0e1a802

└──▷ ALSO FROM THESE RELEASES

. . . .

└──▷ BREAKING ON UPGRADE

! Python 3.10 is no longer supported; Python 3.11 or newer is required.

- ! The `portfilter` module is removed; CDN/WAF port filtering is now built in and on by default — remove `portfilter` from your presets and use `--allow-cdn-ports` to restore the old behavior.
- ! `module_dirs` now only works as a top-level preset key, and BBOT prompts before loading each custom module folder (use `-y` to accept).
- ! Preset conditions are now read-only — they can warn or abort but can no longer modify the preset.
- ! `dnsdumpster` now requires an API key.
- ! `waf_bypass` now only tests scan targets by default.
- ! CPE-based technologies are now named by product, and in-scope events are no longer tagged `affiliate`.
- ! API requests now send a `BBOT/<version>` User-Agent header.
- ! helpers.response_to_json() is removed; use response_to_event_dict() instead.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Non-Human & Agent Identity

Aembit

Feedback ⓘ changelog · Nov 14, 2026 CHANGELOG ↗ RANK

Aembit is a runtime IAM for AI agents and workloads. Enforce access policies, keep downstream credentials out of agents, audit access, and revoke it with a kill switch.

> Aembit added embedded user authorization for MCP servers via OAuth PKCE, new Prometheus metrics covering MCP gateway credential injection and workload event batching, and clearer OIDC client ID visibility.

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

01 New Prometheus metrics for MCP gateway and event batching

NEW 85

`aembit_mcp_gateway_upstream_mcp_server_credential_injections_total` counts requests to MCP servers by whether MCP Identity Gateway added the credential the Access Policy configures; `aembit_mcp_gateway_upstream_mcp_server_http_requests_total` now reports `outcome="auth_error"` when an MCP server returns an HTTP authentication error; `aembit_event_batch_size` is a new histogram recording how many workload events each batch carried for every batch that reached Aembit Cloud. Batch size on a self-hosted MCP Identity Gateway can be capped with `AEMBIT_EVENT_MAX_SEND_BATCH`, e.g. `AEMBIT_EVENT_MAX_SEND_BATCH=100`.

Cap workload event batch size on a self-hosted MCP Identity Gateway to reduce per-flush payload size in high-traffic environments.

env

```
AEMBIT_EVENT_MAX_SEND_BATCH=100
```

— Three metric names plus env var and example value given changelog-20261114-e6762036

02 User authorization into MCP servers via OAuth PKCE NEW

75

Authorize a person into a specific MCP server by naming it in the `mcp_server_url` parameter of a standard OAuth authorization code flow with PKCE — Aembit shows no login page, Identity Provider picker, or consent screen of its own.

— Mechanism and parameter named, but no setup walkthrough changelog-20261114-e6762036

THINNER COVERAGE BELOW

03 Aembit Client ID shown for OIDC Identity Providers IMPROVED

40

A saved OIDC 1.0 Identity Provider now displays the Aembit Client ID that your application sends to Aembit as `client_id`.

— Simple UI display change, no further mechanism given changelog-20261114-e6762036

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ping Identity

Feedback ⓘ changelog · Nov 14, 2026 CHANGELOG ↗

RANK

Ping Identity is an identity and access management platform providing authentication, single sign-on, multi-factor authentication, and identity governance services.

> Ping Identity's rapid-channel update adds named push profiles with new v2.0 journey nodes for multi-brand MFA, a new control to disable `may_act` enforcement for token exchange, and SAML AuthnContext Class Ref prefixing control, alongside smaller performance and usability fixes to journey nodes and the registration flow.

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

01 Named push profiles for multi-brand MFA journeys NEW

80

The Push Notification service now supports named push profiles, letting multiple brands or mobile applications share a single realm. New v2.0 Push Sender, Push Registration, Combined MFA Registration, and Recovery Code Collector Decision journey nodes let you select which push profile a journey uses, with a Legacy Configuration profile available to preserve previous realm-level settings during migration.

When running multiple branded mobile apps in one realm, create a named push profile for each brand so journeys can target the correct push sender.

> 📍 In the admin console, go to Realms › <your-realm> › Services › Push Notification, create a new named push profile, then add a v2.0 Push Sender node to your journey and select that profile.

— Names all four nodes and migration profile but only UI steps, no API.

changelog-20261114-dfc8a178

THINNER COVERAGE BELOW

02 Disable may_act enforcement for token exchange NEW

55

A new control lets you disable `may_act` enforcement for token exchange at the realm or client level, allowing subject tokens from third-party identity providers or authorization servers to be accepted.

— Names the field but gives no exact config path or command. changelog-20261114-dfc8a178

03 SAML AuthnContext Class Ref prefixing control NEW

55

A new `Prefix AuthnContext Class Ref` configuration option on hosted SAML service providers controls whether custom AuthnContext Class Ref values are automatically prefixed, so SPs can send non-URN namespaced references unchanged to the IdP.

When a custom SAML AuthnContext Class Ref already uses a non-URN namespace, disable automatic URN prefixing so the reference reaches the IdP unchanged.

> 📍 In the admin console, go to Realms › <your-realm> › Applications › <hosted-SP> › SAML 2.0, open Advanced settings, and disable `Prefix AuthnContext Class Ref`.

— Exact config option named with UI navigation steps included.

changelog-20261114-dfc8a178

---

04 **Reduced data requests in Identify Existing User and Query Filter Decision nodes**

IMPROVED

50

Version 2.0 of the `Identify Existing User` node and version 2.0 of the `Query Filter Decision` node each reduce the data requested (to identify a user and to evaluate a query filter, respectively), improving journey performance.

— Names both node versions but no steps to use them. changelog-20261114-dfc8a178

---

05 **Hyphenated identifiers in PingOne product connections**

IMPROVED

40

PingOne product connections now accept identifiers containing hyphens, so hyphenated PingOne environment names can be connected and deleted without validation errors.

— Clear before/after but no config key or command given. changelog-20261114-dfc8a178

---

06 **Visual password requirement indicators in registration journey**

IMPROVED

30

The default registration journey now shows password requirement status using checkmarks and Xs rather than color alone.

— Describes the UI change only, no path to configure it. changelog-20261114-dfc8a178

---

07 **Updated PingOne Signals SDK data collection** IMPROVED

20

The PingOne Signals (Protect) SDK now provides updated client-side data collection used by PingOne Protect.

— No detail on what data or mechanism changed. changelog-20261114-dfc8a178

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ changelog · Nov 14, 2026 CHANGELOG ↗

All the latest releases from the 1Password team

1Password: Manage 1Password Environments from the CLI with the new `op environment` command, supporting create, get, list, edit, and delete operations.

└──▷ **TRY IT**

List all 1Password Environments in a vault to audit what environments exist across your org.

```
$ op environment list
```

Retrieve a specific secret from a 1Password Environment for use in a CI pipeline or script.

```
$ op environment secret get <secret-name> --environment <environment-name>
```

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-13) BLOG ↗

RANK

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

> JumpCloud's November 2026 release adds audit-ready Access Certification campaigns, an AI-agent delegated access model, SIEM-ready Directory Insights log export, and user-level Windows policy controls, alongside incremental upgrades to Workflows, Device Lock, App Catalog, SCIM provisioning and browser discovery.

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

01 Access Certification campaigns with audit-ready exports NEW

93

JumpCloud adds an Access Certification feature supporting one-time or recurring review campaigns for SSO applications and directory user groups, multi-stage approval workflows, automated access revocation or identity suspension, and a centralized admin dashboard tracking pending reviews. Completed campaigns can be exported as tamper-proof snapshots in CSV, JSON, or PDF formats with SHA-256 document hashing and an administrator sign-off flow for compliance audits.

Export a frozen, audit-ready Access Certification campaign snapshot in PDF with SHA-256 hashing for a compliance audit.

> 📍 In the Admin Portal, go to Access Certification › open the completed campaign › select Export › choose PDF format › complete the administrator sign-off flow to download the tamper-proof snapshot.

— Full export steps, hashing mechanism and formats specified

Release Notes 2026 (2026-11-13)

02 Agent Delegated Access for AI agents NEW

85

JumpCloud adds Agent Delegated Access (Private Preview), letting an AI agent authenticate with its own JumpCloud-managed identity while acting on behalf of an end user, using token exchange and allowed-audiences configuration, with direct API access without routing calls through AI Gateway.

Enable Agent Delegated Access for an AI agent so it can act on behalf of a user — configure a trusted OIDC application and allowed audiences from the agent's Delegated access tab in the Admin Portal.

📍 In the Admin Portal, open the agent application › select the Delegated access tab › enable delegation › select a trusted OIDC application › configure Allowed audiences for the resources the agent may request tokens for › save.

— Exact Admin Portal steps to enable delegation given Release Notes 2026 (2026-11-13)

03 **SIEM integration for Directory Insights** NEW 73

JumpCloud adds SIEM integration in Directory Insights for continuous identity-event ingestion using an existing API key, supporting RAW and OCSF-normalized formats with JSON or CSV delivery, plus asynchronous historical log export jobs covering up to 7-day windows across a 13-month retention lookback with email notification on completion.

— Formats and retention windows named but no setup command

Release Notes 2026 (2026-11-13)

04 **Device Lock enhancements for Apple MDM** IMPROVED 70

Device Lock for Apple MDM now supports automatic PIN generation and escrow for macOS, retrievable from the Admin Portal under Highlights > Device Lock > View PIN, an optional Lock Screen message and phone number for Mac, iPhone, and iPad, and a macOS Lock Status column in the Devices list.

— Exact Admin Portal navigation to retrieve PIN given Release Notes 2026 (2026-11-13)

05 **User-level Windows registry and MDM policies** NEW 65

JumpCloud adds a User Level Windows Custom Registry Keys Policy that applies settings under `HKEY_CURRENT_USER` when an assigned user signs in, without requiring MDM enrollment or a device restart, and a User Level Custom Windows MDM (OMA-URI) Policy for configuring custom MDM settings scoped to a user's profile on MDM-enrolled Windows devices.

— Named registry path and OMA-URI but no exact config steps

Release Notes 2026 (2026-11-13)

06 **OAuth 2.0 authentication for Workflow connectors** NEW 65

Workflow connectors now support OAuth 2.0 authentication, including Client Credentials and Authorization Code with PKCE, with credentials stored in Secrets Manager and tokens refreshed automatically.

— Auth flows and storage named but no setup walkthrough Release Notes 2026 (2026-11-13)

07 **MCP-enabled AI client management of Workflows** NEW 60

JumpCloud Workflows can now be created, modified, duplicated, and troubleshot through a connected MCP-enabled AI client using natural language, including inspecting

run history and manually triggering workflows.

:
:
:
:

— Capabilities listed but no example command shown Release Notes 2026 (2026-11-13)

THINNER COVERAGE BELOW

08 **Per-application group name override in SCIM** NEW 55

Admins can set a per-application group name override using Show In App As on the User Groups and Agent Groups tab in SCIM provisioning, without changing the underlying JumpCloud group name; the override takes priority over Group Export Attribute Mapping.

— UI location and precedence named, no further mechanism Release Notes 2026 (2026-11-13)

09 **Expanded Workflows template library** NEW 55

The Workflows template library now includes templates for auto-suspending users on inactivity, blocking newly discovered generative AI apps, revoking Mac admin privileges on enrollment, sensitive software install email alerts, user group membership emails, pre-hire holding group placement, and setting device custom attributes on agent enrollment.

— Templates enumerated but usage path not detailed Release Notes 2026 (2026-11-13)

10 **Capability-based app setup for SSO, SCIM, MCP** NEW 55

Capability-based app setup now lets admins select SSO, SCIM, and/or MCP independently when adding an application, and create MCP-only applications available through AI Gateway without configuring SSO.

— Named capabilities but no step-by-step setup Release Notes 2026 (2026-11-13)

11 **Browser discovery via JumpCloud Go extension** NEW 55

Browser discovery in AI & SaaS Management can now start from a JumpCloud-managed device sign-in or via the JumpCloud Go extension deployed through a third-party MDM, using a Client ID and Client secret from Settings.

— Named credentials and deployment path but no full steps Release Notes 2026 (2026-11-13)

12 **App Catalog update notification and self-service controls** IMPROVED 55

App Catalog now lets admins disable end-user update notifications per application with Disable end user update notifications, and lets users initiate required-app updates from the Self Service App Catalog.

— Named setting and surface but limited mechanism detail Release Notes 2026 (2026-11-13)

WAS THIS USEFUL? 👍 👎

Yubico Feedback ⓘ

YubiNation 2026 takeaways: Customers and partners shaping the future of digital trust (2(

BLOG ↗ RANK

Yubico makes YubiKey hardware security keys that provide strong two-factor and passwordless authentication to protect accounts and sensitive systems.

> Yubico introduced two early-access offerings at YubiNation 2026: an identity-lifecycle trust solution called YubiID and a post-quantum-ready YubiKey.

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

01 **PQC-Ready YubiKey public beta** NEW 35

Yubico launched a `PQC-Ready YubiKey` (Public Beta), hardware engineered to operationalize post-quantum defenses without disrupting existing day-to-day workflows.

— Names the product and beta status but lacks technical detail

YubiNation 2026 takeaways: Customers an…

02 **YubiID identity-lifecycle trust solution** NEW 33

Yubico added `YubiID` (Limited Early Access), a solution that carries verifiable trust across the complete identity lifecycle — from initial onboarding to account recovery.

— Describes purpose but no mechanism or access details

YubiNation 2026 takeaways: Customers an…

WAS THIS USEFUL? 👍 👎

Okta Identity Engine

Feedback ⓘ

Version 2026.46.0 (2026-11-11) BLOG ↗ RANK

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

> Okta Identity Engine shipped generally available passkey (FIDO2 WebAuthn) support for user identification policies and expanded AI agent configuration across identification, platform mapping, and custom claims.

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

01 **Passkey support for user identification policies** NEW

40

Passkey (FIDO2 WebAuthn) support for user identification policies is now generally available in Production.

— Names GA status and standard but no setup steps Version 2026.46.0 (2026-11-11)

02 **AI agent configuration expanded** IMPROVED

38

AI agents now support a `source` property for identification and configuration, multiple platform mappings with external ID filtering, and AI agent attributes are now available in custom claims (GA in preview).

— Names config surfaces but no mechanism or usage path Version 2026.46.0 (2026-11-11)

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Commvault Cloud

Feedback ⓘ changelog · Nov 13, 2026 NOTES ↗

BLOG ↗ RANK

Commvault Cloud is a backup, recovery, and cyber-resilience platform for hybrid, multi-cloud, and SaaS workloads.

> Commvault Cloud expanded Protection Groups to cover Amazon Redshift, DynamoDB, and multi-account Azure/GCP onboarding, added SaaS protection for ServiceNow and Zscaler, and introduced AWS European Sovereign Cloud support, Redshift Serverless backups, RDS for SQL Server transaction log backups, and Threat Scan findings triage.

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

01 **Multi-account onboarding in Azure and GCP Protection Groups**

NEW 75

Azure Protection Groups and GCP Protection Groups now support onboarding multiple Azure subscriptions or GCP projects in a single workflow, with rule-based discovery and automatic protection for Azure Virtual Machines, Blob Storage, Azure Files, Azure Data Lake Storage (ADLS), CloudSQL, BigQuery, Bigtable, Spanner, and AlloyDB.

— Enumerates every supported service but no exact setup steps 2026-11-13-ffc0a6fb

02 **Amazon Redshift Serverless snapshot and export backups** NEW

75

Amazon Redshift Serverless data and metadata can now be protected with snapshot and export-based backups, with snapshots supporting cross-region recovery and exported data stored in Commvault AirGap on Amazon S3.

— Names backup types, storage target and recovery mechanism 2026-11-13-ffc0a6fb

03 **SaaS protection for ServiceNow and Zscaler** NEW 60

Commvault now protects and recovers ServiceNow apps and records for organizations that rely on ServiceNow for IT service management, and protects Zscaler security policies, access controls, and configurations, enabling recovery from accidental changes or configuration errors.

— Names both platforms and protected objects but no setup steps 2026-11-13-ffc0a6fb

THINNER COVERAGE BELOW

04 **Redshift and DynamoDB support in AWS Protection Groups** NEW

55

AWS Protection Groups now cover Amazon Redshift and DynamoDB, letting you onboard and protect analytics and application databases across multiple AWS accounts with rule-based policies.

– Names services and mechanism but no exact UI path or command 2026-11-13-ffc0a6fb

05 **Cloud Application Recovery native workflows with expanded Azure coverage**

IMPROVED 55

Cloud Application Recovery (formerly Cloud Rewind) now uses native Commvault protection and recovery workflows with 3x Azure resource coverage, including support for Azure Government.

– States scope and multiplier but no concrete mechanism detail 2026-11-13-ffc0a6fb

06 **Transaction log backups for RDS for SQL Server** NEW

55

Commvault now supports transaction log (t-log) backups for Amazon RDS for SQL Server, enabling point-in-time database restore.

– Specific service and outcome but no exact command shown 2026-11-13-ffc0a6fb

07 **Threat Scan findings triage and failed-scan visibility**

IMPROVED 55

Threat Scan now lets users mark individual or multiple findings as safe to remove non-actionable items from the dashboard while preserving original detections in the index, and surfaces files that failed to scan.

– Describes UI action and effect but no exact navigation path 2026-11-13-ffc0a6fb

08 **AWS European Sovereign Cloud support for database workloads**

NEW 45

Amazon cloud database workloads can now be protected and recovered in AWS European Sovereign Cloud (ESC), supporting European data residency and sovereignty requirements.

– Names the region but lacks configuration specifics 2026-11-13-ffc0a6fb

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

Veeam Kasten

Feedback ⓘ 9.0.11 NOTES ↗ RANK

Kasten K10 is a veeam Kasten is a Kubernetes-native platform for backing up, restoring, and migrating containerized applications and their persistent data.

> Kasten K10 9.0.11 adds HTTP(S) proxy support for outbound connections with automatic OpenShift CA extraction, enriches the location profile picker, and tightens security around the multi-cluster and console plugin proxies with two breaking changes.

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

01 HTTP(S) proxy support with OpenShift CA auto-extraction NEW

85

Kasten K10 now supports HTTP(S) proxies for outbound connections. On OpenShift, setting `global.proxy.openshift.caCertsAutoExtraction=true` automatically extracts the cluster-wide proxy's trustedCA bundle and cluster CA certificates into Kasten's trust bundle, so TLS-intercepting proxies work without manual certificate configuration.

Enable automatic CA extraction on OpenShift so a TLS-intercepting proxy is trusted without manually distributing certificates.

yaml

```yaml
global:
  proxy:
    openshift:
      caCertsAutoExtraction: true
```

— Exact config key and value given, directly runnable 9.0.11

02 Multi-cluster proxy no longer forwards session credentials

60

BREAKING

The multi-cluster proxy no longer forwards the caller's session and authentication credentials to secondary clusters, no longer relays cookie-changing response headers from them to the browser, and WebSocket and other protocol upgrades are no longer supported through the proxy.

— Explains behavior change but gives no migration step 9.0.11

THINNER COVERAGE BELOW

03 OpenShift console plugin proxy drops plain-HTTP port 8080

50

BREAKING

The OpenShift console plugin proxy's plain-HTTP listener on port 8080 is no longer reachable from the openshift-console namespace.

— Names exact port and namespace, no remediation given 9.0.11

04 **Richer location profile selection dropdown** IMPROVED

40

The location profile selection dropdown now shows each profile's provider type, target, region, and immutability status inline.

— Describes UI change but no exact navigation path 9.0.11

└──▷ **BREAKING ON UPGRADE**

! The multi-cluster proxy no longer forwards the caller's session and authentication credentials to secondary clusters, no longer relays cookie-changing response headers from them to the browser, and WebSocket and other protocol upgrades are no longer supported through the proxy.

! The OpenShift console plugin proxy's plain-HTTP listener on port 8080 is no longer reachable from the openshift-console namespace.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ

Tailscale Terraform Provider v0.30.0 (2026-11-13) BLOG ↗ RANK

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

> Tailscale's Terraform Provider gained a new attribute for controlling tailnet route selection, including enabling MagicRoute.

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

01 Route selection control in Terraform provider NEW

75

The `tailscale_tailnet_settings` resource now supports a `route_selection` attribute to set the route selection method for a tailnet, including enabling MagicRoute, e.g. `route_selection = "magicroute"`.

Enable MagicRoute for your tailnet by setting the new route_selection attribute in Terraform.

hcl

```
resource "tailscale_tailnet_settings" "main" {
  route_selection = "magicroute"
}
```

— Names exact resource, attribute, and example value.

Tailscale Terraform Provider v0.30.0 (2…

WAS THIS USEFUL? 👍 👎

Palo Alto Prisma Access

Feedback ⓘ

changelog · Nov 11, 2026 | BLOG ↗ | RANK

Palo Alto Prisma Access is a cloud-delivered security platform providing secure remote access, network security, and zero-trust protection for users and applications globally.

> Palo Alto Networks launched Sovereign Prisma SASE, a regulated-enterprise variant of Prisma SASE adding EU data residency, multi-cloud failover, on-premises processing options, customer-held encryption keys, and compliance-oriented audit logging.

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

01 **Hybrid SASE Private Location for on-prem processing** NEW 55

Sovereign Prisma SASE adds a hybrid architecture option called SASE Private Location, letting organisations process sensitive traffic on-premises and hold their own encryption keys while still feeding anonymised telemetry to global threat intelligence.

– Names a specific product component but no setup detail 2026-11-11-15987456

02 **EU data residency boundary for Sovereign SASE** NEW 48

Sovereign Prisma SASE establishes an EU-region data boundary that restricts administrative access to EU-only personnel and performs in-jurisdiction TLS break-and-inspect, so all traffic processing stays within European jurisdiction.

– Explains mechanism and scope but no config surface named 2026-11-11-15987456

03 **Multi-cloud sovereign backbone with failover** NEW 45

Sovereign Prisma SASE adds a multi-cloud backbone that uses multiple sovereign cloud service providers within the same jurisdiction, so traffic automatically fails over to an alternate provider without leaving the data-residency boundary if one cloud goes down.

– Describes failover mechanism, no concrete configuration steps given 2026-11-11-15987456

04 **Hold Your Own Key encryption model** NEW 43

Sovereign Prisma SASE introduces a Hold Your Own Key encryption model giving organisations direct cryptographic control over their own keys, without hairpinning traffic through external HSMs.

– Named feature with mechanism but no activation steps 2026-11-11-15987456

05 **Audit logs with geo-location data** NEW 30

Sovereign Prisma SASE adds readily available audit trail logs that include user geo-location data to support compliance team reporting.

— Brief description, no log format or access path given `2026-11-11-15987456`

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v4.3.0 NOTES ↗

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 200+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

CISO Assistant adds OIDC-federated service accounts, OAuth-secured MCP, an evidence-review MCP tool, TPRM workspace task filing, and configurable framework defaults.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.3.0 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v4.3.0
```

Add service accounts federated with an external OIDC identity provider for non-interactive, machine-to-machine authentication.

Secure the MCP server with OAuth sign-in so that MCP clients must authenticate before accessing any tool.

Add an evidence-review tool to the MCP server, letting AI agents inspect and act on compliance evidence.

File tasks raised on a third-party audit directly in that third party's workspace via the TPRM module.

Surface extra settings for framework default behaviour, giving administrators control over per-framework defaults.

+3 more

NEW COMPLIANCE FRAMEWORK Hungarian NIS2 compliance preset

UPDATED COMPLIANCE FRAMEWORK NIS2 2024/2690 framework — requirement names added

NEW COMPLIANCE FRAMEWORK ISO 14001:2015 environmental management systems outline (English and French)

WAS THIS USEFUL? 👍 👎

OneTrust

Feedback ⓘ

November 13, 2026 - 202611.1.0 (Fall Release) (2026-11-13) BLOG ↗ RANK

The OneTrust Developer Portal allows you to leverage APIs and SDKs to integrate external systems seamlessly and streamline the data flow with the OneTrust platform.

> OneTrust's Fall 2026 release adds a broad set of new Developer Portal APIs — covering DROP management for data brokers, document exceptions, mobile app consent, and role discovery — alongside enhancements to cookie scanning, consent templates, and script publishing.

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

01 **Per-domain script file generation control** BREAKING

83

Passing the new `useDistinctScriptSrc` request parameter in the Publish Script to Website API generates a unique `otSDKStub.js` per domain and changes the test script source URL. This is a breaking change: any `otSDKStub.js` already implemented on a website must be updated to the new URL after enabling this parameter.

— Names exact parameter, file, and migration step required.

November 13, 2026 - 202611.1.0 (Fall Re…

02 **Mobile App Consent APIs with Web SDK support** NEW 78

Six new Consent Management Platform APIs retrieve and manage mobile app consent experiences: Get Banner Data API, Get Preference Center Data API, Get IAB and Google Vendors API, Log Consent API, Get Universal Consent Purposes API, and Submit Global Opt-Out Request API. Web SDK support was added across four of these — Log Consent API, Get IAB and Google Vendors API, Get Banner Data API, and Get Preference Center Data API — to deliver consistent consent experiences across platforms.

— Names ten API surfaces but gives no endpoint paths or payload shape.

November 13, 2026 - 202611.1.0 (Fall Re…

03 **Document exceptions search APIs in Tech Risk & Compliance**

NEW 75

Search and retrieve document exceptions across Tech Risk & Compliance with the new Search Exceptions API and Get Exception API, now accessible via the `EXCEPTION` and `EXCEPTION_READ` OAuth scopes configured in Global Settings.

— Names APIs and exact OAuth scopes plus where to configure them.

November 13, 2026 - 202611.1.0 (Fall Re…

04 **Cookie scan removal detection and IABC template types**

IMPROVED

75

The Get List of Added or Removed Cookies API now returns a new `isNewlyRemoved` response parameter to detect newly removed cookies in scan results. The Get Template Details API gains support for IAB Canada (IABC) display group types via six new `type` enum values: `IABC_PURPOSE`, `IABC_STACK`, `IABC_FEATURE`, `IABC_SPL_PURPOSE`, `IABC_SPL_FEATURE`, and `IABC_DATA_CATEGORY`.

— Names two APIs, a new field, and six enum values exactly.

November 13, 2026 - 202611.1.0 (Fall Re…

05 **DROP Management APIs for data brokers** NEW

71

Three new APIs let California-registered data brokers manage permanent deletion and opt-out requests: Create Data Subject Request for DROP Record API, Update DROP Record API, and Get List of DROP Records API.

— Names three APIs and their purpose but no method/path or payload detail.

November 13, 2026 - 202611.1.0 (Fall Re…

06 **Roles discovery APIs** NEW

67

Discover and validate tenant roles programmatically with the new Get List of Roles API and Get Role API, accessible via the new `ROLE_READ` OAuth scope.

— Names two APIs and the scope needed to call them.

November 13, 2026 - 202611.1.0 (Fall Re…

THINNER COVERAGE BELOW

07 **Add options to selection-type data elements** NEW

53

The new Add Data Element Option API lets you add new options to selection-type data elements without affecting existing data subjects or historical consent receipts.

— Describes behaviour and named API but no further mechanism.

November 13, 2026 - 202611.1.0 (Fall Re…

08 **High-availability collection point token retrieval** NEW

49

The new Get Collection Point Token API reliably retrieves collection point tokens during scheduled maintenance windows.

— Names one API with a single behavioural claim and no detail.

November 13, 2026 - 202611.1.0 (Fall Re…

└──▷ **BREAKING ON UPGRADE**

! Enabling `useDistinctScriptSrc` in the Publish Script to Website API changes the test script source URL — any `otSDKStub.js` already implemented on the website must be updated to

the new URL.

WAS THIS USEFUL?  

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 12th, 2026 NOTES ↗ RANK

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

> Compyl's latest release adds a new Coupa integration for supplier and approval-chain data, expands Netskope and AWS sync coverage, lets incidents link to multiple third parties, extends custom field search/filter to more grids, and locks down internal workflow-engine APIs behind authentication.

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

01 Coupa integration for supplier and approval data NEW

85

A new Coupa integration syncs `supplier requests`, `custom fields`, `contacts`, `approval chains`, `users`, `lookup values`, and `suppliers` via OAuth, configured with your `instance URL`, `client ID`, `client secret`, and an optional `form filter`.

Connect Compyl to your Coupa instance so supplier data, approval chains, and users sync automatically.

> 📍 In Compyl, go to Integrations › Coupa and enter your instance URL, client ID, client secret, and (optionally) a form filter, then authorize via OAuth and save.

— Names exact config fields and includes setup steps. Nov 12th, 2026

02 Authentication required on workflow-engine APIs BREAKING

60

Internal workflow-engine API endpoints now require authentication and return `401` to unauthenticated callers.

— Names the response code but not specific endpoint paths. Nov 12th, 2026

THINNER COVERAGE BELOW

03 Expanded sync coverage for Netskope and AWS integrations

IMPROVED

55

The Netskope integration now collects DLP incidents, SCIM users, `groups`, and `group memberships`. The AWS integration now syncs `GuardDuty detector`

`status`, `threat intelligence sets`, and IP sets.

— Names synced data types but no setup path given. Nov 12th, 2026

04 **Searchable custom fields on Controls and Clients grids**

IMPROVED 55

Custom fields on the Controls and Clients grids can now be searched and filtered, matching the behavior already available on Risks, Vendors, and other record types.

— Names affected grids but no exact filter syntax. Nov 12th, 2026

05 **Multiple third-party links on incidents** NEW 30

You can now link an incident to multiple third parties directly from the incident record.

— Brief statement with no detail on mechanism or UI. Nov 12th, 2026

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 5.1.0 (2026-11-13) 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 5.1.0 release adds an `aiwatch enforce-test` verification flow, OpenCode session monitoring with Protect/Enforce coverage, policy-aware agent connection UI, and broader redaction and error-masking across AI Watch — alongside a breaking change requiring `RESEND_API_KEY` for self-hosted managed email.

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

01 RESEND_API_KEY required for managed email delivery

BREAKING

95

Managed email delivery (Let Runlayer handle email delivery) for self-hosted deployments now reads exclusively from `RESEND_API_KEY`, with no fallback to WorkOS Vault for its key. Set it in the provisioned tenant secret on ECS or in the `runlayer-app` Secret on Kubernetes before upgrading, or organizations with managed delivery will fall back to custom SMTP or send no email at all, including security-violation alerts.

— Exact env var and secret locations with upgrade consequence spelled out.

5.1.0 (2026-11-13)

02 Enforce test detections via aiwatch enforce-test NEW

85

Run `aiwatch enforce-test` on a device to print a test MCP server entry and two shell prompts that Enforce and Protect block, then confirm each block in the agent reply and in Shadow AI Activity. A remote MCP entry pointed at the Runlayer-hosted test server lets teams verify Enforce blocks remote MCP calls on devices that have no `enforce-test` command yet.

After a policy rollout, verify that Enforce and Protect are blocking on a device by printing the test prompts and checking results in Shadow AI Activity.

```
$ aiwatch enforce-test
```

— Runnable CLI command with exact verification steps and location. 5.1.0 (2026-11-13)

03 OpenCode session monitoring in AI Watch NEW

70

AI Watch now records OpenCode sessions and applies Protect and Enforce to OpenCode tool calls on macOS and Windows, including MCP tools, via a Runlayer-installed OpenCode plugin (OpenCode 1.x, 1.18.29 or newer recommended; OpenCode 2 not yet supported).

— Names versions and platforms but no install command given. 5.1.0 (2026-11-13)

04 **Policy-aware agent connection UI** NEW 60

Agent account connection cards now show deny policies that reach the agent, noting how many tools a policy denies and marking the connector Blocked by policy for blanket denies; denied tools appear first in the tool picker and cannot be selected.

— UI behavior described but no exact navigation path. 5.1.0 (2026-11-13)

05 **Error masking for local MCP server errors** IMPROVED 60

Error messages returned by local MCP servers run through the Runlayer CLI are now masked when PII detection or token masking is set to `mask`, preventing personal data and credentials in errors from reaching the AI client.

— Names the `mask` setting but not its exact config key. 5.1.0 (2026-11-13)

THINNER COVERAGE BELOW

06 **Redaction in AI Watch events and audit exports** IMPROVED 50

AI Watch skill discovery and risk audit events, and audit log exports, now show `***REDACTED***` in place of sign-in details, query values, and recognized tokens in a skill's git remote URL.

— Names the redaction marker but no control to configure it. 5.1.0 (2026-11-13)

└──▷ **BREAKING ON UPGRADE**

! Managed email delivery (Let Runlayer handle email delivery) for self-hosted deployments no longer falls back to WorkOS Vault for its key — `RESEND_API_KEY` must be set in the ECS tenant secret or the `runlayer-app` Kubernetes Secret before upgrading, or organizations with managed delivery will fall back to custom SMTP or send no email at all, including security-violation alerts.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Runtime Defense & Guardrails

doberman-core

Feedback ⓘ v0.20.0 NOTES ↗ RANK

> doberman-core adds AUTH-prompt integration for Claude Code, Codex and Cursor in `doberman dash`, a sandbox 2FA check in `doberman doctor`, single-use TOTP enforcement, a dedicated judge API key variable, and a module-invocation fallback.

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

01 **Judge API key via `DOBERMAN_JUDGE_API_KEY`** NEW 75

The optional judge now prefers the `DOBERMAN_JUDGE_API_KEY` environment variable, falling back to `ANTHROPIC_API_KEY` if unset, letting the judge use a credential scoped separately from the shared Anthropic key.

Point the optional judge at a dedicated API key instead of the shared Anthropic key so credentials stay scoped.

```
$ export DOBERMAN_JUDGE_API_KEY=sk-ant-...yourkey...
doberman dash
```

— Names both env vars and gives a runnable example v0.20.0

02 **Module invocation via `python -m doberman`** NEW 65

doberman-core can now be run without the CLI binary on PATH using `python -m doberman`, for example `python -m doberman doctor`, useful in restricted environments such as CI containers where the binary isn't installed.

Invoke doberman when the `doberman` binary is missing from PATH, such as inside a restricted CI container.

```
$ python -m doberman doctor
```

— Runnable command given, but limited context beyond that v0.20.0

THINNER COVERAGE BELOW

03 **`doberman doctor` sandbox 2FA check** NEW 45

`doberman doctor` now checks whether Claude Code's sandbox keeps your 2FA secret isolated from the agent.

— Names the check but not its mechanism or output v0.20.0

04 **Single-use TOTP codes** IMPROVED 45

TOTP codes are now single-use: a reused or older code is refused even while it is still inside its validity window.

— Explains behaviour change but no config surface to act on v0.20.0

---

05 **`doberman dash` AUTH-prompt integration for Claude Code, Codex, Cursor**

NEW 35

`doberman dash` adds AUTH-prompt integration for Claude Code, Codex and Cursor.

— Only named in the summary, no mechanism or usage detail v0.20.0

WAS THIS USEFUL? 👍 👎

◆ Model & Agent Protection

ToolHive

Feedback ⓘ v0.65.0 NOTES ↗ RANK

> ToolHive adds workload-identity-based client authentication via SPIFFE X.509-SVIDs, cross-authorization-server token delegation through ID-JAG token minting, and trace correlation for audit events.

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

01 **SPIFFE X.509-SVID client authentication** NEW 45

ToolHive can now authenticate clients using SPIFFE X.509-SVIDs, enabling workload-identity-based access without relying on bearer tokens.

– Describes the mechanism but no config keys or flags given v0.65.0

02 **ID-JAG token minting and actor-token validation** NEW 45

ToolHive now mints ID-JAG tokens and validates actor tokens for delegate clients, supporting cross-authorization-server token delegation workflows.

– Explains workflow but lacks concrete API or flag names v0.65.0

03 **Trace and span IDs on audit events** IMPROVED 30

ToolHive now records trace and span IDs on audit events, correlating audit log entries with distributed traces.

– States the change but no schema or field names given v0.65.0

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal

Feedback ⓘ

release-dd7d8cf2 NOTES ↗ RANK

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

> Minimal added a tamper-evident audit trail for boxes with parent-chain verification surfaced in `min doctor`, plus several smaller CLI additions for inspecting and managing boxes.

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

01 **JSON/YAML output for box details** NEW 65

Box details, including its execs, can now be output as JSON or YAML with `min box show -o json` or `min box show -o yaml`.

– Exact flag and formats given; no deeper mechanism explained release-dd7d8cf2

02 **Strip terminal control codes from box logs** IMPROVED 60

`min box logs` now strips terminal control codes when stdout is not a terminal.

– Clear before/after condition and command named release-dd7d8cf2

THINNER COVERAGE BELOW

03 **List and show box types** NEW 40

New `min box` list and show commands let users list and show box types.

– Commands named but behaviour barely described release-dd7d8cf2

04 **Detach keys for exec attach** IMPROVED 30

Users can detach from an exec attach session using the detach keys via `min attach`.

– Thin description with no detail on key bindings release-dd7d8cf2

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Microsoft code-oss

Feedback ⓘ CVE-2026-102230 CVE ↗

STABLE

The code-oss project is an open-source code editor combining editing, debugging, and extensibility for developers across Windows, macOS, and Linux.

`CVE-2026-102230`: Information Disclosure via Integrated Browser File Uploads (CVSS 7.4, affects `< 1.145.3`)

▾ 6 more releases in this issue · 2026-11-10

> CVE-2026-121662 CVE ↗ STABLE
>
> `CVE-2026-121662`: Workspace Trust Bypass via Prototype Pollution in Workspace Configuration (CVSS 8.6, affects `< 1.145.3`)

> CVE-2026-121660 CVE ↗ STABLE
>
> `CVE-2026-121660`: Security Feature Bypass via Inherited Publisher Names in Allowed Extensions Policy (CVSS 7.5, affects `< 1.145.3`)

> CVE-2026-121664 CVE ↗ STABLE
>
> `CVE-2026-121664`: Remote Code Execution via Extension Text Rendered as Markdown (affects `< 1.145.3`)

> CVE-2026-119975 CVE ↗ STABLE
>
> `CVE-2026-119975`: Command Injection via Crafted Publisher Names in the Extension Trust Dialog (affects `< 1.145.3`)

> CVE-2026-120782 CVE ↗ STABLE
>
> `CVE-2026-120782`: Command Execution via Agent-Originated Local Commands (CVSS 6.8, affects `< 1.145.3`)

> CVE-2026-120780 CVE ↗ STABLE
>
> `CVE-2026-120780`: Remote Code Execution via serve-web Release Downgrade (CVSS 7.2, affects `< 1.145.3`)

WAS THIS USEFUL?  

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Nov 14, 2026 CHANGELOG ↗

RANK

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

> Cloudflare's November update adds spending limits and alert thresholds for Workers and Durable Objects, expands D1 and AI Search integration (including a breaking namespace change in the Terraform provider), extends RealtimeKit with token validity windows and noise cancellation, and open-sources its Kitesurf browser engine for AI agents.

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

01 Custom RealtimeKit token validity windows and revocation NEW

90

Set custom validity windows for RealtimeKit auth tokens by passing `not_before` and `expires_on` in the `token` object when adding a participant via `POST /client/v4/accounts/$ACCOUNT_ID/realtime/kit/$APP_ID/meetings/$MEETING_ID/participants`; individual tokens can now be revoked without deactivating the whole meeting.

Issue a RealtimeKit participant token that is valid only during a scheduled session window, so a shared link stops working after the session ends.

```
$ curl
"https://api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/realtime/kit/$APP_ID/meetings/$MEETING_ID/participants" \
  -X POST \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --data '{
    "name": "Mary Sue",
    "preset_name": "<preset_name>",
    "custom_participant_id": "<uuid_of_the_user_in_your_system>",
    "token": {
      "not_before": "2026-12-01T15:00:00Z",
      "expires_on": "2026-12-01T17:00:00Z"
    }
  }'
```

— Full endpoint, field names and runnable example provided. changelog-20261114-79878607

02 **AI Search namespace support and breaking Terraform changes**

90

BREAKING

Terraform provider v5.29.0 adds a `namespace` argument to the `cloudflare_ai_search_instance` resource, enabling management of AI Search instances in any namespace, and introduces `cloudflare_ai_search_namespace` as a new resource. Breaking: `cloudflare_ai_search_instance` import IDs now require the namespace segment (`<account_id>/<namespace>/<id>` , use `default` for the default namespace), the `cloudflare_ai_search_instance` and `cloudflare_ai_search_instances` data sources now require `namespace` , the `filter.namespace` argument on the `cloudflare_ai_search_instance` data source is removed, `cloudflare_ai_search_instances` now returns instances from one namespace only, and changing `namespace` on `cloudflare_ai_search_instance` replaces the instance.

— Exact resource names, argument names and import ID format given.

changelog-20261114-79878607

---

03 **Spending limits for Workers and Durable Objects** NEW

80

Set spending limits and email alert thresholds for Workers and Durable Objects via the Cloudflare dashboard under Spend limits , or with the `cf` CLI using `billing spend-policies` with `--product-id` , `--threshold-cents` , and `--action usage_block` .

— CLI flags and dashboard path named, no single full command shown.

changelog-20261114-79878607

---

04 **Noise cancellation for RealtimeKit web apps** NEW

70

Add noise cancellation to RealtimeKit web apps with the new `@cloudflare/realtimekit-noise-cancellation` package, calling RealtimeKitNoiseCancellation.init({ meeting }) after meeting initialisation. Works in React, Angular, and plain JavaScript on Chrome, Edge, Firefox, and Safari, including mobile.

— Exact package and init call named but no full runnable sample.

changelog-20261114-79878607

---

05 **Local testing for Browser Run quickAction** IMPROVED

70

Browser Run .quickAction() calls can now be tested locally under `npx wrangler dev` , where the browser binding runs against a local Chrome instance instead of requiring remote mode.

— Names method and command, minimal further mechanism detail. changelog-20261114-79878607

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `cloudflare_ai_search_instance` import ID now requires the namespace segment: `<account_id>/<namespace>/<id>`. For instances in the default namespace, use `<account_id>/default/<id>`.

! The `cloudflare_ai_search_instance` and `cloudflare_ai_search_instances` data sources now require a `namespace` argument; set `namespace = "default"` for instances in the default namespace.

! The `filter.namespace` argument on `cloudflare_ai_search_instance` data source is removed; `cloudflare_ai_search_instances` now returns instances from one namespace only.

! Changing the `namespace` argument on `cloudflare_ai_search_instance` replaces the instance.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 14, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue