THE CYBER TOOLCHAIN

NO. 130

# Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 4, 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 204 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 51 of 52 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.

PROTECT

vim ? help

**Get an authenticated scan running against a login flow you could never script**

AI Assisted Authentication drives a real browser through complex multi-step logins and writes out a reusable Zest replay script, so the scan stays logged in on subsequent runs. Authenticated coverage has historically died on hand-crafted auth scripts that break on every UI change; OAuth2 support closes another common gap.

OWASP ZAP

PROTECT

**Fix one base image instead of chasing the same CVE across every downstream container**

Phoenix Security's lineage tracking attributes inherited findings to the base image and names the minimum base version that clears them, turning a long list of duplicate container CVEs into a single bump. StepSecurity adds a suppression-and-resolution workflow for malicious file findings, so a reviewed false positive stops re-firing across Dev Machine Guard and GitHub.

Phoenix Security · StepSecurity

DETECT

**Recover behaviour from a sample that stalled instead of writing it off as benign**

Adaptive Detonation diagnoses why execution halted — missing dependency, environment check, user interaction — and re-runs under corrected conditions. Analysts previously had to spot a dead run themselves and hand-tune the environment, and anti-analysis samples routinely scored clean on the first pass.

Joe Security (Joe Sandbox)

GOVERN

**Run an access review campaign and prove it happened**

JumpCloud's Access Certification puts scheduled reviewer workflows and tamper-evident reports in the directory itself, replacing the spreadsheet-and-email cycle most teams use for audit evidence. Nudge Security adds browser-based discovery of shadow AI agents on top of API integrations, catching tools nobody connected through an admin console.

JumpCloud · Nudge Security

Does **Opus 5** read better? 👍 👎

## OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

# xalgorix

Feedback ⓘ v4.6.177 NOTES ↗ RANK

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

> Xalgorix tightened its scan-completion lifecycle with a bounded verification window and new convergence settings, closed several findings-integrity gaps (proven confirmation marking, verifier provenance, severity caps, case-sensitive dedup), and added Runcrate as a built-in LLM provider plus end-of-scan cost reporting.

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

01 **Bounded verification window and scan-completion lifecycle**

NEW 98

Adds `XALGORIX_CONVERGE_FINISH_IDLE` to control the no-finding tail length before a scan finalizes (default raised from 150 to 500 iterations) when `XALGORIX_CONVERGE_FINISH` is on, and `XALGORIX_CONVERGE_FINISH_LEAD_WINDOW` to give the agent a bounded window (default 100 iterations) to verify high-confidence unverified leads before the scan closes. Scans closed while work remained unexecuted now record `completion: partial` with `stop_reason: untested_at_finish`, list unrun tasks as `untested_at_finish` instead of `not applicable`, and report how much of the plan and route surface was tested. Scans that exhaust or saturate their plan now let the no-new-finding tail decide termination rather than the stuck-loop limit, out-of-band and cross-identity signals stay open for the model to conclude instead of being auto-closed as exhausted, and rate-limit/quota/provider-error retries no longer consume idle limits or the verification window.

Tune scan endurance on hardened targets where findings are rare — raise the idle tail and give the agent more time to verify promising leads before the scan closes.

```
$ XALGORIX_CONVERGE_FINISH=true
XALGORIX_CONVERGE_FINISH_IDLE=1000
XALGORIX_CONVERGE_FINISH_LEAD_WINDOW=200 xalgorix --web
```

Persist tuned convergence settings across restarts by writing them to `~/.xalgorix.env` so every scan inherits them without re-exporting each time.

```ini
XALGORIX_CONVERGE_FINISH=true
XALGORIX_CONVERGE_FINISH_IDLE=800
XALGORIX_CONVERGE_FINISH_LEAD_WINDOW=150
```

— Names every flag and field with runnable example commands. v4.6.177 v4.6.175

02 **Arjun rate-limit flag mapping** IMPROVED 75

The request-rate rewriter now gives arjun its real `-t` and `--rate-limit` options, maps commonly guessed flags (`--threads`, `--chunk`, `--delay`, `--stable-sleep`, `--timeout`, `-oJ`) to their real equivalents, and limits rewriting to commands that actually invoke the tool.

— All flags named but rewriting happens automatically, not manually invoked. v4.6.175

03 **End-of-scan LLM cost reporting** NEW 65

Xalgorix now sums LLM costs reported by the provider (such as OpenRouter's `usage.cost`) across the scan and every subagent, and shows the total in the CLI's end-of-scan usage note.

— Names the field and display location but no config step. v4.6.176

THINNER COVERAGE BELOW

04 **Proven confirmation marking for SQLi, template injection and XXE**

IMPROVED 55

The SQL injection, template injection and in-band XXE verifiers now mark confirmations as proven, including when a confirmation merges into an existing ledger entry, so a scan can no longer end with a confirmed injection left unreported.

— Internal verifier fix with no user-facing flag. v4.6.177

05 **Severity caps applied after CVSS scoring** IMPROVED 55

Per-type severity caps now apply after CVSS scoring, so a model-supplied vector can no longer raise capped finding types such as clickjacking or missing headers to medium or high.

— Named finding types but internal scoring order, not user-actionable. v4.6.177

06 **Case-sensitive ledger and dedup keys** IMPROVED 55

Ledger and finding-dedup keys now preserve path case, lowercasing only scheme and host, so a finding on `/Api/...` is no longer merged with one on `/api/...`, keeping case-variant authentication bypasses as separate findings.

— Concrete before/after example but no user-facing control. v4.6.177

07 **Runcrate added as built-in LLM provider** NEW 55

Xalgorix adds Runcrate, a chat-completions API for open-weight models, to the built-in provider catalog, with model discovery via `/v1/models`.

— Endpoint named but no setup or configuration steps given. v4.6.176

08 **Verifier-engine provenance recorded in ledger evidence**

IMPROVED 50

Ledger evidence now records which engine verifier wrote it, so model-supplied text such as not confirmed is no longer treated as verifier proof.

— Describes mechanism but nothing a user configures. v4.6.177

WAS THIS USEFUL?  

## APPLICATION SECURITY

◆ SAST / SCA / Code Security

# GitHub

Feedback ⓘ

BYOK and enterprise custom models in the Copilot usage API (2026-11-04) BLOG ↗

RANK

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

> GitHub's biggest shipments this window were deeper Copilot usage API visibility into BYOK and enterprise custom models, Dependabot version updates for Agent Package Manager dependencies, a new REST endpoint for cost center lookups, and short-lived auto-refreshing authentication tokens in GitHub Desktop.

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

## 01 BYOK and custom model fields in Copilot usage API NEW

95

The Copilot usage metrics API now exposes bring-your-own-key and enterprise custom model usage in dedicated fields: `used_local_byok` flags per-user local BYOK activity; `totals_by_local_byok` breaks down session, request, prompt, and token counts by `byok_kind` (`local_managed`, `local_user`, `local_auto`, or `remote_user`) and surface (`copilot_cli` or `copilot_app`); `daily_active_local_byok_users`, `weekly_active_local_byok_users`, and `monthly_active_local_byok_users` count distinct BYOK users over time; and a boolean `is_custom_enterprise_model` field on each `totals_by_client[].totals_by_model[]` entry flags enterprise custom model requests across all clients including IDEs. All four fields appear in both enterprise and organization reports, at user and aggregate level, in 1-day and 28-day reports.

Pull a 1-day enterprise usage report and check which users made at least one local BYOK request — useful for auditing BYOK policy compliance.

```
$ curl -H 'Authorization: Bearer <token>' \
  'https://api.github.com/enterprises/<enterprise>/copilot/usage?days=1' \
  | jq '[.[] | select(.used_local_byok == true) | {login, used_local_byok, totals_by_local_byok}]'
```

Separate enterprise custom model token usage from Copilot-hosted model usage across all clients to understand true provider spend.

```
$ curl -H 'Authorization: Bearer <token>' \
  'https://api.github.com/enterprises/<enterprise>/copilot/usage?
days=28' \
  | jq '[.[] | .totals_by_client[] | .totals_by_model[] |
select(.is_custom_enterprise_model == true)]'
```

— Exact field names, values and working curl/jq queries given

BYOK and enterprise custom models in th…

---

02 **Dependabot version updates for Agent Package Manager** NEW

90

Adding an `apm` entry to `.github/dependabot.yml` — specifying the directory containing `apm.yml` and an update schedule — makes Dependabot open pull requests updating git dependencies declared in `apm.yml` and regenerate `apm.lock.yaml` when present. This covers AI agent skills, prompts, and instructions, and is available on github.com for all GitHub plans and in GitHub Enterprise Server 3.24.

Keep AI agent skills and prompts current by having Dependabot open PRs for outdated APM git dependencies on a weekly schedule.

```yaml
version: 2
updates:
  - package-ecosystem: "apm"
    directory: "/"
    schedule:
      interval: "weekly"
```

— Exact config key, files, and a runnable dependabot.yml example

Dependabot version updates support Agen…

---

03 **Short-lived authentication tokens in GitHub Desktop** NEW

60

GitHub Desktop now issues short-lived authentication tokens that renew automatically for both GitHub API requests and Git operations, so no manual re-authentication is needed during normal use. The feature is rolling out gradually on github.com and GitHub Enterprise Cloud (not GitHub Enterprise Server); existing sessions stay signed in after upgrading to 3.6.10, but existing tokens are not auto-migrated — new short-lived tokens are only issued at the next sign-in once the rollout reaches a given installation.

— Explains rollout and migration behavior but no user action required

GitHub Desktop adds short-lived authent…

THINNER COVERAGE BELOW

04 **Cost center lookup REST API endpoint** NEW 55

A new REST API endpoint lets callers look up which cost center a user, organization, repository, or enterprise team belongs to by passing the resource type and name; the response includes the cost center's ID and name, returning an empty `memberships` array when no cost center matches.

— Behavior and response shape given but no endpoint path or method

Look up a resource's cost center with t…

NEW INTEGRATION GitHub Copilot integration in JetBrains Air

NEW SUPPORTED PLATFORM Cloud task execution via Air Teams in isolated Docker containers

NEW SUPPORTED PLATFORM Python 3.15 syntax support

NEW SUPPORTED PLATFORM C23 bit-precise integer support

UPDATED DETECTION BlazingMQ flow summaries for C/C++

UPDATED DETECTION Grape::API remote flow sources for Ruby

UPDATED DETECTION Python XXE, XML-bomb, and NoSQL injection queries respect threat model and data extensions

UPDATED DETECTION Java SSRF and URL-redirection queries recognize case-insensitive URI host comparisons

NEW SUPPORTED PLATFORM Python 3.15 support for Dependabot

└──▷ **BREAKING ON UPGRADE**

! Existing sessions remain signed in after upgrading to 3.6.10, but existing tokens are not automatically migrated to short-lived tokens — new short-lived tokens are only issued at the next sign-in once the rollout reaches your installation.

> GitHub Copilot is now available in JetBrains Air (2026-11-03) BLOG ↗
>
> STABLE
>
> GitHub Copilot now works inside JetBrains Air via ACP, supporting parallel agent sessions, in-context review, and cloud tasks billed to your Copilot plan.
>
> Connect GitHub Copilot to JetBrains Air via the Agent Client Protocol (ACP) to delegate coding tasks, track parallel agent sessions, and review changes in-context inside your JetBrains IDE.
>
> Run Copilot cloud tasks in isolated Docker containers on fresh virtual machines through Air Teams, billed to your existing Copilot plan rather than JetBrains AI credits.
>
> Link your GitHub account for cloud tasks from the browser by navigating to Settings > VCS Providers in JetBrains Air, then selecting a repository, remote branch, and

GitHub Copilot as the agent.

### CodeQL 2.27.3 adds Python 3.15 features and JSON data extensions (2026-11-03)

BLOG ↗ STABLE

CodeQL 2.27.3 adds JSON syntax for data extension files and an opt-in AI threat model for LLM-generated data sources, with C# NuGet proxy/cert passthrough and a changed GitHub Actions CFG library.

Data extension files now accept JSON syntax (`.json` extension) using the same structure as YAML, including Java-style `//` and `/* ... */` comments.

Enable an opt-in `ai` threat model for LLM-generated data sources, letting you define your own sources via data extensions without any built-in source models.

When private registries are configured, CodeQL now passes proxy and certificate environment variables to the NuGet CLI for C# projects.

└─▷ **BREAKING ON UPGRADE**

- ! The GitHub Actions control flow graph (CFG) now uses the shared CFG library; custom queries that rely on specific CFG nodes, edges, or basic block boundaries may need updates.
- ! The legacy `Completion` classes have been removed from the GitHub Actions CFG; use CFG edge labels `DirectSuccessor`, `BooleanSuccessor`, and `ReturnSuccessor` instead.

### Dependabot now supports Python 3.15 (2026-11-03)

BLOG ↗ STABLE

---

WAS THIS USEFUL? 👍 👎

# Phoenix Security

Feedback ⓘ

Container Lineage – Base Image Associations & Remediation (2026-11-03) BLOG ↗

RANK

Phoenix Security correlates scanner, cloud and runtime findings with exploitability and ownership, and uses AI agents to prioritise and open remediation pull requests.

> Phoenix Security shipped container lineage tracking that traces inherited vulnerabilities from base images to downstream containers and surfaces the minimum base version needed to remediate them.

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

## 01 Container lineage for base image vulnerabilities NEW

90

In `Assets > Container`, selecting `Base image` next to `Group by` groups container assets by their shared base image, showing downstream image counts, risk level, risk magnitude and finding counts per severity. From any image or image group, a lineage graph shows the connected base image, downstream image version, image group, registry and associated running containers as a node cluster. In `Vulnerabilities > Remedies > Containers`, grouping by `Image cluster` separates inherited from self-introduced findings by labeling them `SELF` or `BASE IMAGE`, indicating whether a downstream package change or a base-image rebuild is needed, and surfaces the minimum corrected base-image version recommended for remediation.

Identify how many downstream images inherit vulnerabilities from the same base, then find the minimum base version needed to remediate them.

📍 1. Go to Assets › Container. 2. Next to `Group by`, select `Base image`. 3. Expand the affected base-image group to see downstream image count, risk level and finding counts by severity. 4. Open the lineage graph for a downstream image to confirm the base-image digest in use. 5. Go to Vulnerabilities › Remedies › Containers, set `Group by` to `Image cluster`, expand the image group and version, then expand the `BASE IMAGE` remedy to read the minimum corrected base-image version recommended for the rebuild.

– Names exact UI paths, grouping options and labels with worked example

Container Lineage – Base Image Associat…

WAS THIS USEFUL?  

◆ Software Supply Chain Security

# rearm

Feedback ⓘ 26.11.17 NOTES ↗ RANK

ReARM - Release Governance Platform

> ReARM's biggest release this window brings TEA v1.0.0-compliant Transparency Exchange API publishing, signed SBOM exports with licence/supplier data, a new `REGULATORY_DOCUMENT` artifact type for FDA submissions, and agent task boards reaching general availability.

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

## 01 Transparency Exchange API conforms to TEA v1.0.0 IMPROVED

80

rearm's Transparency Exchange API now conforms to the TEA v1.0.0 specification, with a configurable publishing module supporting multiple profiles (Pro) or a single profile (CE), optional TEI stamping on publish, and TEA serving on a ReARM Domain or a separate mapped domain (Pro only).

— Specifies spec version, edition differences, and publishing config options 26.11.17

## 02 Agent task boards reach general availability IMPROVED

73

Agent task boards are now GA for both Pro and CE editions, adding lane paging, board task export as JSON, PR unlinking from the task page, delivery unit display per PR, a list of unlinked PRs, and a threat model specification type.

— Enumerates six concrete sub-features including an export format 26.11.17

## 03 REGULATORY_DOCUMENT artifact type for FDA submissions NEW

65

Adds a new `REGULATORY_DOCUMENT` artifact type, with regulatory features toggled on or off in settings and an option to publish FDA submissions (Pro only).

— Names the artifact type and settings toggle, edition-gated 26.11.17

THINNER COVERAGE BELOW

## 04 SBOM export and scoring improvements NEW

55

SBOMs can now be exported as a signed artifact that carries the declared licence and supplier, with each merged BOM's `serialNumber` lineage assigned a designated owner, and SBOM scores are now computed by rebom and shown in the UI.

— Names serialNumber lineage field but no export command or UI path No detail on scoring criteria or display location

26.11.17

05 **API Key Groups distinguish keys and agents** NEW 50

Adds API Key Groups logic so API keys and agents are distinguished by label and short ID throughout the UI, including the agent session page and release history.

— Describes mechanism but no config surface to act on 26.11.17

06 **Cached totals on Statistics page** IMPROVED 45

The Statistics page now shows cached totals with the date they were counted, displayed as pending until the first daily count runs.

— Explains behaviour but no exact navigation given 26.11.17

07 **Simplified CLI login approval key list** IMPROVED 40

CLI login approval now shows a single key list with bound keys hidden by default.

— Clear UI behaviour change but no command shown 26.11.17

08 **Configurable NetworkPolicy in Helm chart** NEW 40

Helm chart now includes configurable `NetworkPolicy` settings.

— Names a concrete config key but no values or defaults given 26.11.17

09 **New template action and Auto Integrate preview** NEW 25

Adds a new template action and an Auto Integrate preview under Templates and integration.

— Named only in passing with no mechanism given 26.11.17

WAS THIS USEFUL? 👍 👎

## sbomify Feedback ⓘ v26.11.0 NOTES ↗ RANK

sbomify is a product security artifact hub and a trust center.

> sbomify's November release adds human-review gates for AI agent actions, cross-product release linking, CRA Article 14 reporting with ENISA export, and a cluster of VEX visibility and cap improvements.

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

01 **Human review gates for AI agent actions** NEW 70

Changes an AI agent makes to a published advisory now require human approval before they go live, and agent tokens can queue triage decisions for human review by submitting with `review=true`.

– Names the review=true flag and the approval mechanism v26.11.0

02 **VEX visibility, cap, and review improvements** IMPROVED 70

VEX-suppressed findings are now shown everywhere vulnerabilities are counted, including the Trends tab; the VEX statement cap is raised to 50,000 with a warning displayed when that limit is approached; and the VEX review view now shows which versions a submission applies to.

– Names the Trends tab, the 50,000 cap, and the review view v26.11.0

03 **CRA Article 14 reporting and export** NEW 65

sbomify now prepares, tracks, and records CRA Article 14 reports for submission to ENISA's portal. The manufacturer's keyless Sigstore signature is now preserved on CRA export bundles.

– Names CRA Article 14, ENISA portal, and Sigstore signature preservation v26.11.0

THINNER COVERAGE BELOW

04 **Cross-product release linking** NEW 45

A release can now link to another product's release, with linked releases manageable from the release page.

– Describes behaviour and UI location but little mechanism v26.11.0

05 **Trust center advisory page redesign** IMPROVED 35

The trust center advisory page is redesigned around an 'Am I affected?' workflow.

– Names the new workflow but no mechanism detail v26.11.0

06 **API token type shown in token list** IMPROVED 30

Each API token's type is now shown in the token list.

— Bare UI change with no further detail v26.11.0

07 **Vulnerability IDs linked to sbomify Insights** IMPROVED 30

Vulnerability IDs are now linked to sbomify Insights.

— Thin description with no mechanism or scope v26.11.0

WAS THIS USEFUL? 👍 👎

## sbom-operator

Feedback ⓘ 0.49.0 NOTES ↗ RANK

Catalogue all images of a Kubernetes cluster to multiple targets with Syft

> sbom-operator adds flexible naming controls for SBOM storage and Dependency-Track projects, plus configurable deployment timestamps for Dependency-Track uploads.

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

01 **Global naming template and namespace separation for SBOM storage**

NEW 33

A global naming template now controls how SBOM files and Dependency-Track projects are named across all targets, and namespace separation keeps SBOMs for workloads in different Kubernetes namespaces distinct.

— Describes behaviour but no config key or flag name given 0.49.0

02 **Configurable deployment timestamps in Dependency-Track uploads**

IMPROVED 23

Dependency-Track integration now supports configurable deployment timestamps on uploaded SBOMs.

— Bare mention of capability without config detail 0.49.0

WAS THIS USEFUL? 👍 👎

## StepSecurity

Feedback ⓘ

GitHub Checks: Pass on Non-Public Repositories for PWN Request and Script Injection (202

BLOG ↗ RANK

StepSecurity is a security platform that hardens CI/CD pipelines, such as GitHub Actions, by detecting supply-chain risks, enforcing least-privilege runners, and monitoring workflow tampering.

> StepSecurity shipped a configurable exemption for private/internal repos in GitHub Checks and a full suppression-and-resolution workflow for malicious file findings across Dev Machine Guard and GitHub.

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

### 01 Suppression rules and finding actions for malicious files

NEW 77

Dev Machine Guard gains a Suppression Rules page for fleet-wide rules matching on file path, detection rule, SHA-256 hash, campaign, confidence, or any combination, using `*` as a wildcard, applied to all devices or a selected subset; rules cover current findings and files found on later scans, support previewing matching findings before saving, and can be edited or deleted anytime. A suppression rule can also be created directly from a finding with its detection rule and file path pre-filled, optionally limited to that finding's device. Individual findings can be resolved or suppressed with an optional comment, with a new Reason column showing why; a detection rule behind a finding can be deactivated for every device from the Findings page with a required reason. On GitHub, suppressed findings likewise show their reason in a new Reason column, and new versions of a file inherit the suppressed or resolved status until explicitly changed.

Suppress a known test-fixture binary fleet-wide before it floods every future scan with false positives.

📍 In Dev Machine Guard, go to Suppression Rules › New Rule. Set File path to the fixture path (e.g. `tests/fixtures/malware_sample.exe`), set Detection rule to the matching rule name, leave Device as All devices, choose reason Used in tests, then click Preview matching findings to verify scope before saving.

Quickly scope a suppression to exactly one endpoint when triaging a finding that is unique to a developer's machine.

📍 On the Dev Machine Guard Findings page, open the finding, click Create suppression rule (file path and detection rule are pre-filled), then click Limit to this

device , select reason Accepted risk , add an optional comment, and save.

Suppress a reviewed test fixture in GitHub so every existing and future version of the file is silenced without repeated manual triage.

📍 In StepSecurity GitHub Findings, locate the file finding, click Suppress , select reason Used in tests , add an optional comment, and confirm — the status applies to all versions of the file across every branch, pull request, and tag.

— Names exact match fields and UI flow but no API/config surface

Malicious Files: Suppression Rules and …

02 **Pass on non-public repos for PWN Request and Script Injection checks**

NEW 60

A new Pass on non-public repositories setting, configured independently per control under GitHub Checks > Configuration, lets PWN Request and Script Injection checks pass on private and internal repositories while still failing on public ones.

Reduce alert noise from private and internal repos while keeping strict enforcement on public repos for PWN Request findings.

📍 In the StepSecurity console, go to GitHub Checks › Configuration, select the PWN Request control, and enable Pass on non-public repositories .

— Exact UI path and setting name given, no deeper mechanism

GitHub Checks: Pass on Non-Public Repos…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

## infisical

Feedback ⓘ v0.167.6 NOTES ↗

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

Infisical adds MCP tool allowlists, project soft-delete with restore, IdP MFA trust for SAML/OIDC, and secret-scanning navigation improvements.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.167.6
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.167.6
```

Add an MCP tool allowlist for MCP services in agent-vault, letting operators restrict which tools an MCP service may call.

Add a grace period and restore capability for deleted secret manager projects, so accidental deletions can be recovered.

Trust identity provider MFA for SAML and OIDC logins, so users who already completed MFA at their IdP are not prompted again.

Link resource rows in secret scanning to their filtered findings for faster triage.

Add descriptions to agent-vault proxies.

+2 more

NEW INTEGRATION CockroachDB dynamic secret support via PostgreSQL provider

WAS THIS USEFUL? 👍 👎

# GitGuardian SaaS

Feedback ⓘ announced · Nov 2, 2026 BLOG ↗

RANK

GitGuardian SaaS is a cloud-hosted security platform that detects exposed secrets, API keys, and credentials in source code repositories and developer workflows.

> GitGuardian shipped a native Wiz integration that syncs open secret incidents from version control into Wiz daily as secret findings, letting cloud security teams see leaked secrets alongside cloud risk without leaving Wiz.

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

01 **Wiz integration for secret incident sync** NEW 91

Configure the integration at `Settings > Integrations > Wiz` by pasting the Client ID, Client secret, Authentication URL, and API endpoint URL from Wiz, then map Info and Unknown severities to Wiz equivalents. Once configured, GitGuardian syncs open secret incidents daily from GitHub, GitLab, Bitbucket, and Azure DevOps into Wiz as native secret findings, each carrying repository, file path, commit, detector, validity status, severity, and a link back to the GitGuardian incident. Secret values never leave GitGuardian — Wiz receives only a SHA-256 hash of the value plus metadata — and validity context surfaces live, still-working keys ahead of inactive ones. The daily sync can be paused or resumed at any time from the same `Settings > Integrations > Wiz` page.

Connect GitGuardian to Wiz so your cloud security team sees leaked secrets alongside cloud risk without leaving their tool.

> 📍 In Wiz, open the GitGuardian integration page and copy the Client ID, Client secret, Authentication URL, and API endpoint URL. In GitGuardian, go to Settings › Integrations › Wiz, click Configure, paste the four values, choose how to map Info and Unknown severities to Wiz severities, and click Add integration. Findings appear in Wiz within 24 hours.

— Names exact config path, fields, sources, and security mechanism

launch-20261102-556efbd2

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

# DataDome

Feedback ⓘ React Native SDK for Fetch 3.0.0 (2026-11-03)

BLOG ↗ RANK

DataDome protects web, mobile, and API applications by detecting and blocking bot and fraud traffic.

> DataDome shipped a major rewrite of its React Native SDK for Fetch, moving to a client-based API built on CoreDataDome with a long list of breaking removals and new platform requirements.

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

## 01 createDataDomeFetch client API for React Native NEW

95

The React Native SDK for Fetch 3.0.0 is rewritten on top of CoreDataDome SDK, replacing `DataDomeFetch` and the `DataDome` singleton ( `DataDome.configure` , `DataDome.getInstance` ) with createDataDomeFetch(dataDome), which takes a `DataDome` client created with new DataDome({ clientKey }) and returns a `fetch-compatible` function that automatically adds DataDome headers to every request and checks every response for a challenge. It accepts a `string` , `URL` , or `Request` as fetch input, and a `Headers` object, an array of pairs, or a plain object as headers; it re-exports `DataDome` , `ResponseValidationStatus` , `DataDomeConfiguration` , and `DataDomeCookie` from `@datadome/react-native-core-datadome` , returns the server response as-is if the CoreDataDome SDK fails, and the package itself is rewritten in TypeScript.

Initialize DataDome protection for all fetch calls in a React Native app using the new client-based API.

javascript

```
import { DataDome } from '@datadome/react-native-core-datadome';
import { createDataDomeFetch } from '@datadome/react-native-
datadome-fetch';

const dataDome = new DataDome({ clientKey: '<your-client-key>'
});
const protectedFetch = createDataDomeFetch(dataDome);

// Use protectedFetch anywhere you would use fetch
const response = await
protectedFetch('https://api.example.com/data');
```

— Names exact API, config object and includes runnable example

React Native SDK for Fetch 3.0.0 (2026-…

02 **Breaking changes for SDK 3.0.0 migration** BREAKING 85

The package is renamed from `@datadome/react-native-datadome` to `@datadome/react-native-datadome-fetch` (update imports and package.json); `@datadome/react-native-core-datadome` is now a required peer dependency; `DataDomeFetch`, `DataDome.configure`, `DataDome.getInstance`, `DataDomeModal`, `DataDomeModule`, and `enableVerboseLogs` are removed; the per-request `datadome` option is removed since headers and challenge checks now apply automatically to every request; the `react-native-webview` peer dependency is removed because challenges are now presented natively by the CoreDataDome SDK; and the SDK now requires React Native 0.85 or newer with the New Architecture, iOS 15.1 or newer, and Android API 24 or newer.

— Lists every removed API and new platform requirement verbatim

React Native SDK for Fetch 3.0.0 (2026-…

└──▷ **BREAKING ON UPGRADE**

! The package is renamed from `@datadome/react-native-datadome` to `@datadome/react-native-datadome-fetch`; update all imports and your package.json dependency.

! `@datadome/react-native-core-datadome` is now a required peer dependency and must be installed alongside this package.

! `DataDomeFetch`, `DataDome.configure`, and `DataDome.getInstance` are removed; replace them with createDataDomeFetch(new DataDome({ clientKey })).

! The `datadome` per-request option is removed; the returned fetch function now applies DataDome headers and challenge checks to every request automatically.

! `DataDomeModal`, `DataDomeModule`, and `enableVerboseLogs` are removed.

! The `react-native-webview` peer dependency is removed; challenges are now presented natively by the CoreDataDome SDK.

! React Native 0.85 or newer with the New Architecture is now required, along with iOS 15.1 or newer and Android API 24 or newer.

WAS THIS USEFUL? 👍 👎

## Fastly

Feedback ⓘ Fastly Terraform Provider (beta) 0.2.7 (2026-11-03)

BLOG ↗ RANK

Fastly is an edge cloud platform offering content delivery network, load balancing, edge computing, and web application security services for developers and enterprises.

> Fastly's beta Terraform provider gained new resources for audit log routing and domain management, plus the early building blocks of a legacy-to-auto migration toolchain, while the JavaScript SDK picked up a runtime upgrade.

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

01 **Terraform legacy-to-auto migration tooling** NEW 80

Adds an initial `fastly-tf-migrate` CLI foundation with version reporting (e.g. `fastly-tf-migrate --version`) and dedicated build/test targets for Terraform migration workflows, a typed legacy-to-`_auto` compatibility registry to support migration analysis, and support for legacy simple product enablement state during VCL-to-CDN Auto migration in `fastly_service_cdn_auto`.

Check the version of the new migration CLI before running a Terraform legacy-to-auto migration analysis.

```
$ fastly-tf-migrate --version
```

– Includes a runnable CLI command for version check

Fastly Terraform Provider (beta) 0.2.7 …

02 **Domain management routing config resources** NEW 60

Adds `fastly_routing_config_auto` and `fastly_routing_config_domain_link` Terraform resources for domain management routing configs with automatic versioning.

– Two resources named, no runnable example given Fastly Terraform Provider (beta) 0.2.7 …

THINNER COVERAGE BELOW

03 **Audit log event routing resource** NEW 55

Adds the `fastly_audit_log_event_mapping` Terraform resource and data source to route audit log events to integrations.

– Names resource and data source but no usage example

Fastly Terraform Provider (beta) 0.2.7 …

04 **StarlingMonkey runtime upgraded to v0.4.0** IMPROVED 30

The Fastly JavaScript SDK now uses StarlingMonkey v0.4.0.

– Only a version bump is stated, no change details JavaScript SDK 4.1.0 (2026-11-03)

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

## PortSwigger Burp Suite Feedback ⓘ

DAST 2026.10.4 (2026-11-03) 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 DAST 2026.10.4 adds GitHub issue tracking for findings and a timestamped evidence card view.

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

01 **GitHub issue tracking for DAST findings** NEW 37

Burp Suite DAST can raise and track security issues directly in GitHub from within the tool, with issue evidence shown as timestamped cards in the findings interface.

– Describes the feature but no endpoint, flag or setup steps given.

DAST 2026.10.4 (2026-11-03)

WAS THIS USEFUL? 👍 👎

# OWASP ZAP

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗ RANK

The ZAP by Checkmarx Core project

> OWASP ZAP's headline addition is AI Assisted Authentication, which uses an LLM to drive a real browser through complex multi-step login flows and auto-generate a reusable Zest replay script; alongside it ZAP shipped OAuth2 authentication support, an in-browser Zest Script Viewer, Client Spider improvements, a new passive rule for server-side technology CVEs, and smaller Automation Framework and add-on updates.

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

## 01 AI Assisted Authentication for complex login flows NEW

100

ZAP adds an AI Assisted Authentication method that uses an LLM to drive a real browser through multi-step login flows that Browser Based Authentication heuristics cannot handle, then automatically converts the successful session into a reusable Client Script Authentication Zest script. Configure it in the Automation Framework by setting the authentication method to `ai` with parameters `loginPageUrl`, `browserId`, `loginPageWait`, and optional `hint` under the `authentication` block, or trigger it programmatically via the `aiAssistedAuthentication` action in the authentication API. The LLM provider is configured via the `llm-config` Automation Framework job (added by the LLM Support add-on v0.1.0), specifying `defaultProvider`, `defaultModel`, and a `providers` list with `type`, `apiKey`, `endpoint`, and `models` fields. Each login step — page description sent, actions chosen, and reasoning — can be monitored in the AI Assisted Auth tab inside the LLM Chat tab, and free-text `hint` guidance (e.g. 'Set the domain dropdown to Project1 before clicking Next.') steers the LLM past steps it can't infer from the page alone. It can also be tried interactively via Tools > Authentication Tester ( Ctrl+T ), selecting AI Assisted as the Auth Method.

Run a scan against an app with a complex multi-step login that Browser Based Authentication cannot handle, letting the LLM work out the flow once and generate a reusable script.

```yaml
env:
  contexts:
    - name: "ai-auth-context"
      urls:
        - "https://www.example.com"
      authentication:
        method: "ai"
        parameters:
          loginPageUrl: "https://www.example.com/login"
          browserId: "firefox-headless"
          loginPageWait: 5
          hint: "Set the domain dropdown to Project1 before
clicking Next."
      verification:
        method: "autodetect"
      sessionManagement:
        method: "autodetect"
      users:
        - name: "test-user"
          credentials:
            username: "[email protected]"
            password: "${AUTH_PASSWORD}"
jobs:
  - type: llm-config
    parameters:
      defaultProvider: "my-provider"
      defaultModel: "my-model"
      providers:
        - name: "my-provider"
          type: "OPENAI_COMPATIBLE"
          apiKey: "${LLM_API_KEY}"
          endpoint: "https://llm.example.com/v1"
          models:
            - "my-model"
  - type: "requestor"
    parameters:
      user: "test-user"
      requests:
        - url: "https://www.example.com/"
          method: "GET"
          responseCode: 200
```

Open the Authentication Tester interactively to test a complex login page with AI assistance and optional hints before committing to automation.

📍 1. Ensure the LLM Support add-on is configured with a provider. 2. Open Tools › Authentication Tester (or press Ctrl+T / Cmd+T on macOS). 3. Set Auth Method to `AI Assisted`. 4. Enter the Login URL, Username, and Password. 5. In the Hints field, enter any guidance, e.g. 'Set the domain dropdown to Project1 before clicking Next.' 6. Press Test and watch progress in the results panel and the AI Assisted Auth tab inside the LLM Chat tab.

— Full mechanism, config keys, API action and runnable example given

`launch-20261103-b7ebd607` `launch-20261102-b931fc36`

---

02 **Authentication diagnostics failure retention controls**

IMPROVED 73

ZAP's Authentication Helper add-on (v0.43.0) adds `type` and `count` parameters to the diagnostics automation job, letting it retain only the last failed attempt with `auth_on_failure` or a rolling window of failures with `auth_failure_rolling`.

— Named parameters and values but no worked example `launch-20261102-b931fc36`

---

03 **Client Spider targeting and AJAX Spider fallback** IMPROVED

73

ZAP's Client Side Integration add-on (v0.32.0) adds a `runOnlyIfModern` option to the `spiderClient` Automation Framework job to run the Client Spider only when a `Modern Web Application` alert has been raised, and now stands in for the AJAX Spider's `spiderAjax` automation job and `ajaxSpider` API using the Client Spider when the AJAX Spider add-on is not installed.

Run the Client Spider only on pages ZAP has already flagged as modern web apps, avoiding unnecessary crawl overhead on traditional server-rendered sites.

```yaml
- type: spiderClient
  parameters:
    runOnlyIfModern: true
```

— Named config key plus a runnable YAML example `launch-20261102-b931fc36`

---

04 **Insights job exit threshold by severity level** IMPROVED

70

ZAP's Insights add-on (v0.7.0) adds an `exitAutoLevel` option (accepting `medium` or `high`) to the insights Automation Framework job, preferred over the now-secondary `exitAutoOnHigh` field.

Restrict a CI scan to exit non-zero only when high-severity findings exceed the threshold, using the new exitAutoLevel setting in the Insights Automation Framework

job.

```yaml
- type: insights
  parameters:
    exitAutoLevel: high
```

— Named option with values and a runnable config example launch-20261102-b931fc36

05 **Passive rule for server-side technology CVEs** NEW 60

ZAP's Passive scanner rules (v77) add a Server Side Technology Version with Known Vulnerabilities rule that checks technology versions disclosed in response headers against known vulnerabilities, initially covering PHP versions via php-src security advisories.

— Describes mechanism and initial scope but no usage path launch-20261102-b931fc36

06 **In-browser Zest Script Viewer** NEW 60

ZAP's Browser Extension (v0.3.0, included in Client Side Integration v0.34.0) adds a Script Viewer that lets you see, edit, replay (including step-by-step), load, and save a Zest script without leaving the browser.

— Clear capability list but no exact UI navigation given launch-20261102-b931fc36

THINNER COVERAGE BELOW

07 **OAuth2 authentication method** NEW 55

ZAP's Authentication Helper add-on (v0.43.0) adds an OAuth2 authentication method supporting the `client_credentials` and `password` grant types, configurable via the desktop UI, the Automation Framework, and the API; the Automation Framework (v0.61.0) correspondingly adds `oauth2` authentication method support.

— Names grant types and surfaces but no config example launch-20261102-b931fc36

08 **Network add-on event notifications** NEW 35

ZAP's Network add-on now notifies other add-ons when an outbound proxy is enabled (v0.30.0) and when requests fail to send, such as DNS errors, connection refused, or timeouts (v0.32.0).

— Internal add-on signal with no end-user action point launch-20261102-b931fc36

09 **Automation plan run-time limit** NEW 20

ZAP's Automation Framework (v0.61.0) allows defining how long a plan can run.

— No key name or mechanism given, bare description launch-20261102-b931fc36

WAS THIS USEFUL? 👍 👎

## REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

# Joe Security (Joe Sandbox) Feedback ⓘ

The Sandbox That Fights Back: Introducing Adaptive Detonation (2026-11-03)

BLOG ↗ RANK

Joe security (joe sandbox) is a automated Malware Analysis - Development and Licensing of Automated Malware Analysis Tools to Fight Malware

> Joe Sandbox introduces Adaptive Detonation, an AI agent within Joe Sandbox AI that diagnoses why a sample's execution stalled and intervenes to re-run it under corrected conditions, exposing behavior that a single static run would miss.

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

## 01 Adaptive Detonation AI agent for stalled samples NEW

72

Adaptive Detonation runs a sample, identifies what blocked execution (an evasion check, a missing prerequisite, or a failure), applies a targeted intervention, and re-analyzes the sample under the new conditions to surface behavior the initial run missed. Interventions include supplying missing launch arguments — for example recovering and injecting the required `S1=1` MSI installer property — without modifying the sample package; adjusting the analysis environment, such as setting the guest machine's date to a required window or advancing the guest clock (demonstrated by a two-day advance to bypass a file-age loader delay) before re-running the unchanged sample in Joe Sandbox Cloud; and running preparation code to create required files, directories, or registry values, or installing a missing software dependency such as Python in a location the sample already searches. Analysis reports now include the agent's findings: what blocked execution, the runtime and code evidence, the recommended action, alternative explanations, and remaining uncertainty. It is available as part of Joe Sandbox AI and has been running in Joe Sandbox Cloud Basic.

.
.
.
.

— Mechanism and examples detailed, but feature runs automatically with no user-facing command

The Sandbox That Fights Back: Introduci…

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

# EMBA

Feedback ⓘ v2.0.5-ice-edt NOTES ↗ RANK

EMBA - The firmware security analyzer

> EMBA's latest release adds a component URL field to its JSON config output, extends binary version detection to the `iw` wireless utility, and ensures CycloneDX SBOM output carries valid license entries for identified components.

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

01 **Component URL in JSON config output** IMPROVED 40

EMBA now adds a component URL field to the JSON config output for identified firmware components, giving analysts a direct reference link alongside each detected component.

— Names the JSON config field but no further mechanism given. v2.0.5-ice-edt

02 **Version detection for iw wireless utility** NEW 30

EMBA adds version detection support for the `iw` wireless utility component, expanding its binary component identification coverage.

— Names the specific utility but lacks detail on detection method. v2.0.5-ice-edt

03 **Valid license entries in CycloneDX SBOM output** IMPROVED 30

EMBA's CycloneDX SBOM output now includes valid license entries for identified components.

— Names the SBOM format but gives no mechanism or scope. v2.0.5-ice-edt

WAS THIS USEFUL? 👍 👎

# Finite State Platform

Feedback ⓘ

October 30, 2026 - Component versions, kernel vulnerabilities, and .NET end of life (202

BLOG ↗ RANK

Finite State Platform is a firmware and software supply chain security tool that analyzes binaries and embedded devices to identify vulnerabilities, SBOM components, and risks.

> Finite State Platform added two new detection capabilities in this release: identifying Brotli as a software component on firmware builds, and surfacing end-of-life status for .NET runtime components.

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

01 **.NET runtime end-of-life status** NEW 38

Finite State Platform now shows end-of-life status for .NET runtime components, including NuGet runtime packs and installer names.

— Names runtime packs and installers but no concrete usage steps

October 30, 2026 - Component versions, …

02 **Brotli component detection in firmware** NEW 30

Finite State Platform now identifies Brotli as a component and reads its version on firmware builds that export one.

— Names the component but no mechanism or UI path given

October 30, 2026 - Component versions, …

WAS THIS USEFUL? 👍 👎

## NETWORK SECURITY

◆ Email Security

# Abnormal AI

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗ RANK

Abnormal AI is a behavioral AI security platform that protects email, identity, and AI—and stops insider threats. Trusted by 5,000 organizations.

> Abnormal AI spun AI Governance out into a standalone product that organizations can buy and activate without first owning Inbound Email Security, adding discovery, inventory, chat-risk, and spend/compliance reporting for AI tools and agents across an organization.

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

## 01 AI Governance launches as standalone product NEW 65

AI Governance is now available as a standalone product that customers can purchase and activate independently, without first buying Inbound Email Security, giving it a new, separate entry point into the Abnormal platform. It is activated from the AI App Store and connects to Microsoft 365 via the Graph API or to Google Workspace in read-only mode; other products such as Inbound Email Security, Identity Threat Protection, and Infiltration Prevention can be added later from the same platform.

Activate AI Governance as a new standalone entry point into Abnormal AI's platform without purchasing Inbound Email Security first.

> 📍 In the Abnormal AI App Store, select AI Governance and complete onboarding. Connect to Microsoft 365 (via the Graph API) or Google Workspace in read-only mode. Additional products such as Inbound Email Security, Identity Threat Protection, and Infiltration Prevention can be added from the same platform at any time.

Activate AI Governance for an organization that does not yet use Abnormal for email, to begin discovering unsanctioned AI tools within about a week.

> 📍 In the Abnormal portal, go to the AI App Store and select AI Governance, then follow the activation steps to connect your environment.

— Names connection methods and activation path but no deeper config detail.

launch-20261103-a4f45b17 launch-20261103-aa26d2bf

THINNER COVERAGE BELOW

## 02 AI chat insights for sensitive data exposure NEW 53

AI chat insights surfaces sensitive data — including credentials, personal information, and financial data, as well as content inside uploaded files and images — shared with AI

assistants, with findings ranked by severity and mapped to compliance frameworks.

— Specifies data types and severity ranking but no actionable workflow.

launch-20261103-aa26d2bf launch-20261103-a4f45b17

03 **AI spend and compliance reporting** NEW 53

AI spend and compliance reporting consolidates AI subscriptions and invoice emails — including spend outside approved procurement — surfacing unsanctioned spend alongside IT-approved spend, and maps findings to reference compliance frameworks to support audit and board reporting.

— Names spend sources and reporting purpose but no specific framework or command.

launch-20261103-aa26d2bf launch-20261103-a4f45b17

04 **AI agent inventory and access auditing** NEW 50

AI agent inventory finds AI agents across connected platforms, audits the access each holds, and flags those that are dormant, orphaned, or overprivileged so teams can revoke stale access.

— Describes flagging mechanism but no concrete action path given.

launch-20261103-a4f45b17 launch-20261103-aa26d2bf

05 **AI tool discovery and classification** NEW 38

AI tool discovery identifies every AI tool in use across connected environments and classifies each one as sanctioned, tolerated, or unsanctioned.

— States classification states but no mechanism or UI path. launch-20261103-aa26d2bf launch-20261103-a4f45b17

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

# Sophos Access Points

Feedback ⓘ

1.10.3324NewReleased on November 3, 2026 NOTES ↗ RANK

> Sophos Access Points adds 6 GHz band steering for AP6 420E and 840E devices and improved SKU validation by location, alongside profile management and traffic-view usability updates.

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

01 **6 GHz band steering for SSIDs** NEW 70

Configure 6 GHz band steering for SSIDs on AP6 420E and 840E access points, choosing whether devices are steered to the 5 GHz band first, the 6 GHz band first, or balanced across all frequency bands, with a manual configuration option also available.

— Names exact devices and steering modes but no CLI/API detail

Fusion 2026.44Released on October 12, 2…

THINNER COVERAGE BELOW

02 **Access point profile management improvements** IMPROVED 45

Sophos Fusion profile management now supports sorting and searching access point profiles by name, and the platform now supports up to 50 access point profiles, including the base profile.

— Names the limit and UI capability but no flags or endpoints

Fusion 2026.44Released on October 12, 2…

03 **Top 5 clients graph device navigation** IMPROVED 45

Clicking a device in the Top 5 clients graph now navigates to that device's summary page for bandwidth investigation.

— Clear UI navigation path described Fusion 2026.44Released on October 12, 2…

04 **Traffic category and uncategorized views paginated with bar graphs**

IMPROVED 45

Categories and details and Uncategorized traffic views now show all categories, ports, and domains across multiple pages, with associated pie charts replaced by bar graphs that update based on the visible list.

— Describes the UI change and behaviour but no exact navigation path

Fusion 2026.44Released on October 12, 2…

05 **SKU validation by access point location** IMPROVED 30

Sophos Access Points now validates SKUs based on the access point's physical location, via an updated Sophos Fusion agent.

— States the change but no mechanism or config detail given

1.10.3324NewReleased on November 3, 2026

WAS THIS USEFUL? 👍 👎

# Vectra AI

Feedback ⓘ v9.15 - Control (2026-10-20) BLOG ↗ RANK

This is the home page for docs.vectra.ai.

> Vectra AI extended automatic lockdown with longer duration options to cover extended periods without manual intervention.

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

01 **48- and 72-hour automatic lockdown durations** IMPROVED

60

Automatic lockdown now supports 48- and 72-hour durations for HostID Lockdown, Traffic Lockdown, and Active Directory Accounts, in addition to all existing options, letting analysts contain risk across weekends, holidays, or low-coverage periods without manual intervention.

.
.
.
.

— Names exact durations and affected lockdown types, no UI path given

v9.15 - Control (2026-10-20)

NEW DETECTION Command-and-Control detection expansion

NEW DETECTION Suspicious SMB Activity Coverage

UPDATED DETECTION Kerberoasting Host-Centric Cipher Downgrade detection

UPDATED DETECTION SPA Detections Available by Default

NEW SUPPORTED PLATFORM Azure Esv5 Support for Cloud Brain

UPDATED SUPPORTED PLATFORM 100 Gbps Sensor, Brain, and Stream Support

UPDATED SUPPORTED PLATFORM Encapsulated Endpoints Support for Mixed Mode Appliances

v9.15 - Observe (2026-10-20) BLOG ↗ STABLE

Vectra AI adds post-quantum cryptography metadata in Stream and Recall, a new PQC Dashboard, and expanded JA4L fingerprinting across general TCP traffic.

Vectra AI now captures post-quantum cryptography metadata in Stream and Recall, giving teams visibility into which cryptographic methods are in use across network communications.

Access a new Post-Quantum Cryptography Dashboard under `Dashboard > Post-Quantum Cryptography Dashboard` to monitor and track PQC adoption across your network.

Vectra AI now extends JA4L fingerprinting to cover general TCP traffic, beyond the previous SSL/TLS and HTTPS-only scope, providing richer network context for investigation.

v9.15 - Platform (2026-10-20) BLOG ↗ LATEST STABLE

Vectra AI v9.15 reorganizes platform navigation under Threat, Exposure, Observability, and Platform, and adds a dedicated Exposure section combining Inventory and Findings.

Platform left navigation is reorganized under Threat, Exposure, Observability, and Platform sections; Respond is now Priorities, Investigate is now Event Search, and Discover is now Dashboards.

A new Exposure section in the platform navigation brings together Inventory and Findings in one dedicated area.

└──▷ **BREAKING ON UPGRADE**

- ! Esv3-series Azure instance types can no longer run after Nov 15, 2029; Cloud Brain deployments must migrate to Esv5-series before that date.
- ! Sensors with Match/SPA enabled will be rebooted as part of the v9.15 upgrade.

WAS THIS USEFUL? 👍 👎

◆ Network Analysis & Monitoring

## Security Onion

Feedback ⓘ announced · Nov 2, 2026 BLOG ↗

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

Security Onion 3.4.0 adds Agentic AI Automations for Pro users; Security Onion 2.4 reached EOL on October 1, 2026, requiring migration to Oracle Linux 9.

Security Onion Pro now includes Agentic AI Automations, enabling automated AI-driven workflows within the platform.

UPDATED SUPPORTED PLATFORM Suricata updated to 8.0.8

UPDATED SUPPORTED PLATFORM Zeek updated to 8.0.11

UPDATED SUPPORTED PLATFORM Elasticsearch updated to 9.4.8

UPDATED SUPPORTED PLATFORM ElastAlert2 updated to 2.32.0

└──▷ BREAKING ON UPGRADE

! Security Onion 2.4 reached End of Life on October 1, 2026; Security Onion 3 requires Oracle Linux 9, so installations on other distributions must perform a fresh install.

WAS THIS USEFUL? 👍 👎

## CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

# Lacework FortiCNAPP

Feedback ⓘ updated · Nov 4, 2026

NOTES ↗ RANK

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

> Lacework FortiCNAPP expands its identity risk detection to cover Azure, adding three new categories of exposure and privilege findings.

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

01 **Azure identity risk detection** NEW 40

Lacework FortiCNAPP now detects Azure identity risks across three new categories: resource exposure, privilege passing, and credential exposure.

— Names three categories but no mechanism or scope detail snapshot-20261104

WAS THIS USEFUL? 👍 👎

◆ CNAPP

# Plerion

Feedback ⓘ CDR findings, easier to find and easier to read:

NOTES ↗ RANK

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

> Plerion added asset ownership so teams and people can be assigned, discovered, and tracked per asset, and introduced new AWS Security Incident Response case detections that require a CDR StackSet update to function.

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

01 **Asset ownership assignment and discovery** NEW 53

Assign a team, a person, or both to each asset so owners are notified when action is needed, and filter assets by owner. Owners are discovered automatically from existing tags, with clearly-matching owners merged and kept current as tags change. Ownership rules can be added for assets tags miss, matching by tag, integration, resource type, region or name, and assets nobody owns yet can be viewed.

— Mechanism described but no concrete UI path or config key given Every asset, an owner:

UPDATED INTEGRATION CDR Event tab summaries for AWS services

└──▷ **BREAKING ON UPGRADE**

! Existing CDR StackSet deployments must be updated to the latest template to send AWS Security Incident Response events to Plerion; without this update, `PLERION-CLOUDTRAIL-86` through `PLERION-CLOUDTRAIL-89` will not fire.

> CDR findings, easier to find and easier to read: NOTES ↗ LATEST STABLE
>
> Plerion adds a two-view findings list separating configuration and CDR event findings, plus a narrative Event tab with a JSON toggle on CDR findings.
>
> Plerion's findings list now has two views next to Group by : Configuration findings (opens by default, showing cloud, workload, and code configuration checks) and Event findings (showing CDR findings, newest first).
>
> The Event tab on a CDR finding now shows a plain-language summary of who did what, where, and when, with a JSON toggle to view or copy the raw event.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

# OpenCTI

Feedback ⓘ 7.261103.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI's latest release adds Elasticsearch 9 support, broader license compatibility, custom fields, mass workflow operations, deferred form intake, and expanded ingestion health checks.

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

01 **Elasticsearch 9 support via dual client** NEW 50

OpenCTI now supports Elasticsearch 9 servers through a new dual Elastic client, broadening the range of supported search backends.

— Names the version and mechanism but no migration detail 7.261103.0

02 **Mass operations on workflow status and playbooks** NEW 50

Users can now apply mass operations on workflow status and playbooks across multiple objects at once.

— Describes scope of bulk action but no UI path given 7.261103.0

03 **Deferred object creation in form intake** IMPROVED 45

Form intake now supports deferred creation of objects — objects prepared in a form are not created until the form is submitted.

— Explains the before/after behaviour clearly, no config details 7.261103.0

04 **Expanded ingestion health checks** IMPROVED 45

Ingestion health checks now include health settings and validation of users with incorrect capabilities.

— Names new checks but no thresholds or commands 7.261103.0

05 **Global TLS certificate verification switches** NEW 35

New switches added to platform configuration let administrators control TLS certificate verification globally.

— States the switch exists but not the exact config key 7.261103.0

06 **Filigran license format accepted** NEW 35

OpenCTI now accepts a Filigran license installed as the platform license, in addition to the native OpenCTI license format.

– Names the license types but no setup steps 7.261103.0

07 **Custom fields on entities** NEW 35

A new custom fields feature lets users enable additional custom fields on entities.

– Bare description of the feature without scope or limits 7.261103.0

WAS THIS USEFUL? 👍 👎

## Lookyloo

Feedback ⓘ v1.43.0 NOTES ↗ RANK

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

> Lookyloo adds a new pivoting method based on Cloudflare NS record pairs to find related captures, alongside a breaking change raising the minimum Python version.

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

01 **Cloudflare NS-pair pivoting** NEW 55

Users can pivot on the indexed tuple of Cloudflare NS record pairs to find other captures containing domains with the same suspected owner.

:

– Explains mechanism but no UI path or command shown v1.43.0

02 **Python 3.11 minimum requirement** BREAKING 40

Python 3.11 or later is now required; deployments running Python 3.10 or older will not be supported.

– Clear migration note but no further detail v1.43.0

└──▷ **BREAKING ON UPGRADE**

! Python 3.11 or later is now required; deployments running Python 3.10 or older will not be supported.

WAS THIS USEFUL? 👍 👎

◆ SIEM

# Databricks

Feedback ⓘ

Control Lakebase sink upserts with merge strategies (2026-11-03) BLOG ↗ RANK

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

> Databricks shipped four distinct updates this window: configurable merge strategies for Lakebase sink upserts, a Beta cross-workspace job-completion trigger for Lakeflow Jobs, general availability for Zerobus Ingest into default storage, and expanded Unity Gateway budget tracking for AI Functions.

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

## 01 Cross-workspace upstream job completion triggers (Beta)

NEW INTEGRATION NEW 80

A Lakeflow Job can now be configured to run automatically when one or more upstream jobs complete, including jobs in other workspaces that share the same Unity Catalog metastore, by setting dependencies on the downstream job. Configure via the job's Triggers tab: select `Add trigger`, choose `Job completion`, and specify the upstream job as a dependency.

Trigger a downstream Lakeflow Job in another workspace as soon as an upstream job finishes, without manual scheduling.

> 📍 In the Databricks UI, open the downstream job, go to Triggers, select `Add trigger`, choose `Job completion`, and specify the upstream job (including cross-workspace jobs attached to the same Unity Catalog metastore) as a dependency.

— Includes exact UI navigation steps to configure

Trigger Lakeflow Jobs when upstream job…

## 02 Merge strategies for Lakebase sink upserts NEW 67

Structured Streaming writes to Lakebase now support merge strategies for upserts, letting you keep the first or last value, ignore nulls, or select values by an ordering column during the upsert. Requires Databricks Runtime 19 or above.

— Names mechanism and runtime requirement but no config syntax shown

Control Lakebase sink upserts with merg…

## 03 AI Functions spend tracked in Unity Gateway budgets IMPROVED 60

Unity Gateway budgets now track spend from task-specific AI Functions — `ai_parse_document`, `ai_extract`, and `ai_classify` — across both batch and real-time execution paths.

— Names exact functions and paths but no config or command shown

AI Functions usage is now included in U…

THINNER COVERAGE BELOW

04 **Zerobus Ingest into default storage (GA)** NEW SUPPORTED PLATFORM

IMPROVED 40

Zerobus Ingest in Lakeflow Connect now supports ingesting into tables backed by default storage, reaching general availability.

— Names the surface but no mechanism or usage steps given

Zerobus Ingest into default storage is …

Okta connector is available in Beta (2026-11-03) BLOG ↗ STABLE

Monitor data freshness in Lakeflow Connect database connector pipelines (Public Previe

BLOG ↗ STABLE

Databricks adds a Data freshness column and event-log fields to Lakeflow Connect CDC pipelines, with alerting support, in Public Preview.

A new Data freshness column in the Lakeflow Connect pipeline tables list shows time since the latest ingested source commit for each CDC table.

Data freshness and end-to-end lag are queryable from the ingestion pipeline's event log, with alert creation supported.

WAS THIS USEFUL? 👍 👎

## EXPOSURE & VULNERABILITY MANAGEMENT

◆ Security Validation (BAS & Automated Pentest)

# Horizon3 NodeZero

Feedback ⓘ updated · Nov 4, 2026 NOTES ↗

RANK

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

> NodeZero's October 2026 release adds new pentest scope and throttling controls, entity-type routing in the Integration Router, Git repository reconstruction in External Asset Discovery, and expanded attack surface coverage including SQL Server trusted links, Azure storage contents, and AI-agent-driven WebApp Pentest guardrails.

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

01 **Pentest scan impact controls** NEW 70

A new `Scope > Exclude TCP Ports` option lets operators exclude specific TCP ports from a pentest to protect fragile assets, and a new Nuclei Scan Intensity (throttling) option controls Nuclei scan intensity to reduce load on sensitive targets.

Protect a fragile OT device on a non-default port by excluding it from TCP scanning during a pentest.

> 📍 In the pentest scope configuration, go to Scope › Exclude TCP Ports and enter the port numbers (e.g. 102, 502) you want NodeZero to skip.

Throttle Nuclei scans against a production web app to avoid overwhelming it during active-hours testing.

> 📍 In the pentest configuration, locate the Nuclei Scan Intensity option and set it to a lower throttling level before launching the op.

— Names both config options with UI steps but no exact values/ranges given

snapshot-20261104

THINNER COVERAGE BELOW

02 **Entity-type routing in Integration Router** NEW 50

Integration Router workflows can now be filtered and routed by entity type, letting users choose which entity types (e.g. findings, assets) are sent to each destination workflow.

Send only critical findings to your SIEM integration while routing asset data to a separate ticketing workflow.

📍 In the Integration Router settings, open the target workflow, select `Entity Types`, and choose the specific entity types (e.g. findings, assets) to send to that destination.

— Describes mechanism and UI path but no API/config key named snapshot-20261104

03 **AI-agent WebApp Pentest with guardrails** NEW 45

WebApp Pentest now uses AI agents constrained by guardrails, with new Risk Assessment and Safety Overrides features to identify potentially risky routes that can be blocked or tested.

— Names the new features but gives no mechanism detail or example snapshot-20261104

04 **Git repo reconstruction in External Asset Discovery** NEW 40

External Asset Discovery now reconstructs exposed Git repositories found on the discovered attack surface and captures service banners on nonstandard ports, broadening visibility into exposed assets.

— Explains new discovery behavior but no usage example or config detail snapshot-20261104

05 **Azure storage content testing** IMPROVED 40

Azure tests now cover the contents of blobs and file shares inside Azure storage accounts and respect minimum runtimes during testing.

— Names scope of new coverage but no example of use snapshot-20261104

06 **Expanded target-app login support** IMPROVED 40

Target-app login now supports SSL VPNs and identity-gated portals, and maintains persistent authenticated sessions during testing.

— Names new supported login types but no concrete steps snapshot-20261104

07 **SQL Server trusted link testing** NEW 30

NodeZero now tests Microsoft SQL Server trusted links as part of its attack surface coverage.

— Bare statement of new test coverage, no mechanism or example snapshot-20261104

08 **Filtering and sorting in Engagement Reports** IMPROVED 25

Engagement Reports now support filtering and sorting.

— Bare one-line statement with no detail on fields or options snapshot-20261104

09 **Scope control in subdomain discovery** IMPROVED 20

Subdomain discovery now offers more user control over scope.

— Vague statement with no specifics on what scope controls were added `snapshot-20261104`

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ **Vulnerability Management**

## oxo

Feedback ⓘ `v2.22.3` NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> OXO extended its message protocol with new types for ASN fingerprinting and ASN WHOIS asset scanning.

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

01 **ASN fingerprinting message type** NEW 50

oxo now supports the `v3.fingerprint.asn` message type for ASN fingerprinting between agents.

— Names the message type but no mechanism detail `v2.22.3`

02 **ASN WHOIS scannable asset type** NEW 50

oxo adds the `v3.asset.asn.whois` message type to support ASN WHOIS as a scannable asset.

— Names the message type but no mechanism detail `v2.22.2`

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

## syzkaller

Feedback ⓘ changes since 2026-11-02 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds AI-driven bisect-cause and bisect-fix workflows and per-bug embargo durations to its aflow and dashboard subsystems

syzkaller adds `bisect-cause` and `bisect-fix` workflows to the `pkg/aflow` subsystem to identify, via AI, the commits that introduced or fixed a kernel bug using crash reports, reproducers, and optional root-cause analysis.

syzkaller's aflow patching workflow now supports To/Cc recipient changes during patch iteration, letting reviewers request email address additions or removals alongside tag changes.

syzkaller's dashboard now supports per-bug embargo durations.

WAS THIS USEFUL? 👍 👎

## runZero

Feedback ⓘ runZero 5.1.261103.0 (2026-11-03) BLOG ↗

An asset discovery and inventory platform that identifies and maps all devices on networks without requiring agents.

UPDATED SUPPORTED PLATFORM BACnet protocol attribute collection improvements

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

## ITarian

Feedback ⓘ

Release Notes: ITarian Platform & Agent Update -October 21, 2026 (2026-10-21)

BLOG ↗

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

ITarian adds Windows ARM64 enrollment, customizable 33-column device-list template views, a folder-scoped procedure search, and a new LocalSystem run permission.

Control who can run procedures as the LocalSystem user with the new `configuration.procedures.run_as_localsystem` role permission — users without it see only Run as Logged in User in the Run Procedure dialog.

Limit procedure searches to the current folder with the new Search within current folder only checkbox in the filter panel, on by default when a folder is selected.

NEW SUPPORTED PLATFORM Manage Windows ARM64 devices by selecting Windows ARM64 in Device Enrollment ( Enroll and Protect and Just Enroll ) and Bulk Installation Package, alongside existing x64, x86, and x86&x64 options.

Save, load, update, reset, and delete device-list layouts using the new Template view button in the device list toolbar, choosing from 33 columns grouped as General, Desktop, and Mobile — shared across all admins within the same portal.

Review which users hold a role instantly, as the Assign Users tab now opens on the Assigned filter by default, with All , Assigned , and Unassigned filter options available.

WAS THIS USEFUL? 👍 👎

### IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

# Varonis

Feedback ⓘ updated · Nov 4, 2026 NOTES ↗ RANK

Varonis is the leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis added a Collibra data catalog integration that surfaces sensitive data classification and security context alongside governance metadata.

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

01 **Collibra data catalog integration** NEW 45

Varonis now integrates with Collibra to enrich catalog assets with sensitive data classification insights, security context, ownership, lineage, and policy information directly inside the Collibra data catalog.

·
·

— Names integration and metadata types but no setup steps or config details

snapshot-20261104

WAS THIS USEFUL? 👍 👎

◆ Privileged Access Management

# BeyondTrust

Feedback ⓘ

Remote Support 26.3.1 release notes (2026-11-03) BLOG ↗ RANK

BeyondTrust is a privileged access management company offering security software that controls, monitors, and audits privileged accounts, credentials, and remote access for enterprises.

> BeyondTrust shipped Remote Support 26.3.1 with new Configuration, session-policy, and AI Session Summary APIs reachable through an MCP gateway, mandatory end-user consent for RDP screen/file/shell access, FIDO2 support on Pathfinder, and expanded Jump Client, traffic-routing, and automation scheduling controls, while Base 8.4.0 added coordinated cluster updates, a Canadian French admin interface, and independent burst/sustained rate limiting.

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

## 01 Configuration, session-policy, and AI summary APIs NEW

83

New APIs let automation and AI assistants manage Remote Support without custom REST integrations: the Configuration API can now be accessed through the MCP gateway so AI assistants and automation tools can read and change RS configuration settings, with read and write access controlled at the organization level; attended and unattended session policies can be assigned to a group policy through the API, removing the manual UI step from provisioning scripts; and AI Session Summaries can be retrieved through public APIs for reporting, storage, and analysis.

Let an AI assistant or MCP-compatible automation tool read and update Remote Support configuration without building a custom REST integration.

📍 Steps: In `/login`, go to Management › MCP Gateway. Enable read and/or write access to the Configuration API at the organization level. Point your MCP-capable AI assistant or tool at the MCP gateway endpoint — it can now query and modify RS configuration settings directly.

— Named APIs and gateway with setup steps, but no exact endpoint paths

Remote Support 26.3.1 release notes (20…

## 02 Jump Client lifecycle controls NEW

75

Jump Client maintenance settings in `/login` now let administrators set how many days a Gateway can go without connecting before the Jump Client uninstalls itself

(default remains 30 days), and control Jump Client upgrades by Jump Group so staged rollouts can test a subset of endpoints before upgrading all.

Roll out a Jump Client upgrade to a pilot Jump Group first to validate compatibility before pushing to all endpoints.

> 📍 Steps: In `/login`, go to Jump › Jump Clients. Select a Jump Group and open its upgrade settings. Choose the specific group to upgrade and set the maintenance window. Confirm; remaining Jump Groups are unaffected until you explicitly include them.

— Named settings and default with rollout steps Remote Support 26.3.1 release notes (20…

03 **Traffic-node routing controls** NEW 75

Cloud administrators can control whether the traffic node selector appears on the representative web login page via a new setting on the Cluster tab in `/login` (when hidden, RS connects representatives to the closest available traffic node automatically), and can choose a traffic-node routing mode — SRV Record Lookup (default), Random, or Timezone Offset — under Management > Cluster > Traffic Nodes.

— Exact navigation path and named routing modes Remote Support 26.3.1 release notes (20…

04 **Per-asset session recording override** NEW 73

Session recording can now be turned on or off for a single Asset or Asset Group via Jump Policies settings, overriding the global recording policy; only administrators can change this setting and every change is logged.

Stop session recording for an Asset Group where recording is prohibited by policy, without disabling recording globally.

> 📍 Steps: In `/login`, go to Jump › Jump Policies and open or create a policy linked to the target Asset or Asset Group. Toggle the session recording override to off for that Asset or Asset Group and save. Only administrators can change this setting, and every change is logged.

— Exact steps and governance detail given Remote Support 26.3.1 release notes (20…

05 **Scheduled Endpoint Automation jobs** NEW 70

Endpoint Automation jobs can now be scheduled on a one-time or recurring basis for specific Jump Groups, tags, or endpoints, with configurable run windows, concurrency limits, and automatic retries.

— Mechanism named, no exact config or API path Remote Support 26.3.1 release notes (20…

06 **End-user consent controls in RDP sessions** NEW 65

Remote Support now requires end-user consent before a representative can share the screen, access files, or access the shell in an active RDP session, with consent re-requested whenever the session switches mid-connection; administrators can also set a session policy to start screen sharing in view-only mode, requiring the representative to request end-user approval before taking control.

— Mechanism and scope described but no exact config path

Remote Support 26.3.1 release notes (20…

07 **AI Session Summaries ServiceNow integration** IMPROVED 65

AI Session Summaries now writes summaries into matching ServiceNow incidents automatically and pulls ticket context from ServiceNow, covers the full session recording rather than only the first ten minutes, and is now available for Atlas deployments.

— Clear before/after but no UI or API path given Remote Support 26.3.1 release notes (20…

08 **Coordinated cluster and failover updates** IMPROVED 65

Clustered and failover deployments now update as a unit — the primary node schedules both base and product updates and every node installs them in the same maintenance window, with per-node status visible on the cluster Updates page.

— Mechanism explained with UI location named Base 8.4.0 release notes (2026-11-03)

09 **RADIUS authentication deprecation** DEPRECATED 60

RADIUS authentication will be removed in the 27.1 release; administrators must migrate to SAML, LDAP, or another supported authentication method before upgrading to 27.1.

— Clear deprecation timeline and migration path Remote Support 26.3.1 release notes (20…

THINNER COVERAGE BELOW

10 **Total Active Users count on Status page** NEW 55

The Status page in `/login` now shows a Total Active Users count — users who signed in within the last 365 days and still exist — immediately after the Time Zone field.

— Exact UI location but a single-line addition Remote Support 26.3.1 release notes (20…

11 **Jump Client Comment and Tag in session reports** IMPROVED 55

Session reports in `/login` and the Reporting API now include the Comment and Tag values set on Jump Clients.

— Named fields but a single-line addition Remote Support 26.3.1 release notes (20…

12 **FIDO2 authentication on Pathfinder** NEW 50

Remote Support on Pathfinder now supports FIDO2 authentication, allowing representatives and administrators to register and sign in with FIDO2 security keys or platform authenticators.

— States capability without setup steps Remote Support 26.3.1 release notes (20…

13 **Independent burst and sustained rate limiting** IMPROVED

50

Rate limiting now limits burst and sustained request rates independently, so short legitimate traffic bursts are no longer throttled as though they were sustained abuse.

— Explains before/after but no config knobs named Base 8.4.0 release notes (2026-11-03)

14 **Canadian French admin interface** NEW

25

The appliance administrative interface is now available in Canadian French.

— Single-line language addition, no further detail Base 8.4.0 release notes (2026-11-03)

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! RADIUS authentication will be removed in the 27.1 release; migrate to SAML, LDAP, or another supported authentication method before upgrading to 27.1.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

# Yubico

Feedback ⓘ

YubiNation 2026: Coming together to shape the future of digital trust (2026-11-03)

BLOG ↗ RANK

Yubico makes YubiKey hardware security keys that provide two-factor and passwordless authentication, protecting accounts against phishing and unauthorized access.

> Yubico previewed two new directions: an identity lifecycle service called YubiID and a post-quantum cryptography-ready YubiKey.

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

01 **Post-quantum cryptography-ready YubiKey** NEW 45

Yubico announced a post-quantum cryptography-ready YubiKey with crypto-agile hardware built on open standards, now available in public beta for customer evaluation.

— Names PQC and beta availability but no specs or enrollment steps.

YubiNation 2026: Coming together to sha…

02 **YubiID identity lifecycle service** NEW 37

Yubico introduced YubiID, an identity lifecycle service covering enrollment, credential issuance, authentication, and recovery — including digital wallets and AI workflows. It is now in limited early access for design partners.

— Describes scope but no concrete interface or sign-up path.

YubiNation 2026: Coming together to sha…

WAS THIS USEFUL? 👍 👎

# 1Password

Feedback ⓘ 1Password for Android 8.12.44 (2026-11-03)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password's latest cross-platform release adds post-quantum passkey support on Android alongside a breaking rename of its MCP server tools, plus smaller platform-specific additions like iPhone Duo authentication, Zen Browser auto-connect, and a setting to hide generated passwords.

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

## 01 MCP server tools renamed from variables to secrets

BREAKING

75

The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets`; any automation, integration, or script calling the old tool names will break.

— Exact old and new tool names given, no migration timeline noted

1Password for Linux 8.12.44 (2026-11-03) 1Password for Mac 8.12.44 (2026-11-03) 1Password for Windows 8.12.44 (2026-11-…

## 02 Hide generated passwords setting in browser extension

NEW

75

A new Hide generated passwords toggle in Settings > Security hides generated passwords from the inline menu, the password generator, and the extension popup.

Keep generated passwords out of view on shared or observed screens so autofill suggestions stay hidden until explicitly requested.

> 📍 In the browser extension, go to Settings › Security and enable Hide generated passwords.

— Exact setting path and affected surfaces given, directly actionable

1Password in the browser 8.12.44 (2026-…

## 03 Watchtower 2FA naming and QR scan unified in Android

IMPROVED

65

Watchtower's multi-factor authentication card is now called 2FA available to match the passkeys card naming; the Scan QR code option is renamed Scan 2FA QR code and is now available on any login in the browser extension, including from the Watchtower multi-factor authentication banner; and the Watchtower Use 1Password for 2FA banner

in item details now links directly to instructions for turning on two-factor authentication on that item's website.

— Names exact UI labels and locations but changes are cosmetic/navigational

1Password for Android 8.12.44 (2026-11-…

THINNER COVERAGE BELOW

04 **Post-quantum passkey support on Android** NEW 40

1Password for Android can now save and sign in with post-quantum passkeys on sites that support them.

— States the capability but no mechanism or supported-site list

1Password for Android 8.12.44 (2026-11-…

05 **Zen Browser extension auto-connect on Mac** IMPROVED 40

Setting up the 1Password browser extension in Zen Browser now automatically connects to the 1Password app.

— Brief description with no setup steps or conditions

1Password for Mac 8.12.44 (2026-11-03)

06 **iPhone Duo authentication method on iOS** NEW 35

1Password for iOS now supports iPhone Duo as an authentication method.

— Names the feature but no detail on enabling or scope

1Password for iOS 8.12.44 (2026-11-03)

07 **API Credential item field renamed to key name** IMPROVED 35

The first field in API Credential items is now labeled key name instead of username.

— Cosmetic label change with no further behavioral detail

1Password for iOS 8.12.44 (2026-11-03)

└──▷ BREAKING ON UPGRADE

! The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets` — any automation or integration calling the old names will break.

! The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets` — any automation or script calling the old names will break.

! The 1Password MCP server tools `list_variables` and `append_variables` are renamed to `list_secrets` and `append_secrets` — any automation or scripts calling the old names will break.

WAS THIS USEFUL? 👍 👎

## Clerk Feedback ⓘ

Group webhooks, group reads, and pause for Directory Sync (2026-11-03) BLOG ↗

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

Clerk adds the ability to read a Directory Sync group and its members via the Backend API, and to pause a directory without removing user access.

NEW INTEGRATION Clerk's Backend API now supports reading a group and its members via Directory Sync.

Pause a Directory Sync directory without removing any user access.

WAS THIS USEFUL? 👍 👎

## Nudge Security

Feedback ⓘ

Detect cryptocurrency wallet addresses shared in AI conversations

NOTES ↗

RANK

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

> Nudge Security added browser-based discovery of shadow AI agents to complement its existing API integrations, extending AI agent coverage to 23 platforms, and added governance actions directly in the AI agents inventory.

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

01 **Browser-based AI agent discovery** NEW 78

Nudge Security now discovers shadow AI agents through browser activity in addition to existing API integrations, covering platforms whose APIs are limited or unavailable to security teams. Browser-based discovery covers Cursor automations, Retool, Zapier Agents, ChatGPT workspace agents, OpenAI Agent Workflows, Atlassian Rovo (Jira, Bitbucket), Zoom AI Workflows, HyperAgent, HubSpot agents, Datadog, and Amazon, bringing total platform coverage to 23 across both discovery channels. Users toggle AI agent discovery on in the browser extension settings to auto-populate the Identities > AI agents inventory as employees view or create agents, with no new rollout needed for existing users.

Activate browser-based agent discovery to surface shadow AI agents your employees are building in tools like Cursor or Zapier Agents that have no enterprise API.

> 📍 In the Nudge Security browser extension settings, toggle AI agent discovery on. Agents will start appearing in Identities › AI agents inventory as employees view or create them.

— Names mechanism, 11 platforms, total count, and exact toggle location

launch-20261102-45c5a114

THINNER COVERAGE BELOW

02 **Governance actions in AI agents inventory** NEW 50

Security teams can set approval status, assign a technical contact, and nudge the agent creator to confirm intent or remediate, all directly from the AI agents inventory.

— Names three actions and location but no exact navigation steps launch-20261102-45c5a114

NEW DETECTION Cryptocurrency wallet address detection in AI conversations

NEW INTEGRATION Duo Security connected app

Detect cryptocurrency wallet addresses shared in AI conversations NOTES ↗

STABLE

Strengthen your SaaS security posture with our latest connected app for Duo Security

NOTES ↗ LATEST STABLE

---

WAS THIS USEFUL?  

# RIG Security

Feedback ⓘ

Rig Security + HiBob: Identity Security That Moves with Your People

NOTES ↗

RANK

Rig security is a rig ties every account, token, and AI agent to the human behind it. Map attack paths, enforce least privilege, and detect identity threats in one platform.

> Rig Security shipped a HiBob integration that pulls HR context into its identity graph, powering department-based policy scoping, employee-to-identity correlation, orphan detection, and leave-aware activity context.

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

### 01 HR context enrichment via RICE correlation engine NEW

65

RICE, Rig's correlation engine, now matches each HiBob employee record to known identities in the graph, surfacing HR context — employment status, job title, department, manager, site, and time-off dates — alongside effective access.

— Names engine and exact fields surfaced, but no UI or query path

Rig Security + HiBob: Identity Security…

THINNER COVERAGE BELOW

### 02 HiBob integration via read-only OAuth NEW

52

Rig Security adds a HiBob integration, connected through a single read-only OAuth consent in the Connections Hub under Partnerships, with no agent to deploy.

Connect HiBob to Rig Security so HR data flows into the identity graph for lifecycle-aware access decisions.

> 📍 In the Rig Security console, go to Connections Hub › Partnerships › HiBob and complete the read-only OAuth consent flow.

— Names exact UI path and consent mechanism, but no API or flag detail

Rig Security + HiBob: Identity Security…

### 03 Department-based populations in Gatewatch policies NEW

45

HiBob departments are now available as populations in Gatewatch policies, letting teams set different guardrails for Finance versus Engineering.

— Clear capability with example but no setup steps given

Rig Security + HiBob: Identity Security…

04 **Orphan identity flagging from HiBob mismatch** NEW 45

Identities that cannot be tied to any HiBob employee are now flagged as orphans, with reporting lines showing which team an account belongs to and who to contact.

— Describes behavior and reporting output, no exact navigation given

Rig Security + HiBob: Identity Security…

05 **Leave-status context for account activity** NEW 40

Rig Security now reads leave status from HiBob to add context for activity on accounts belonging to employees currently on approved leave, with detections based on HR context — such as a privileged action from an employee on approved leave — in development.

— Core detection capability explicitly noted as still in development

Rig Security + HiBob: Identity Security…

WAS THIS USEFUL?

## authentik

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗ RANK

The authentication glue you need.

> authentik's 2026.11 release introduces post-quantum cryptography support across the platform, adding ML-DSA certificates, OAuth2 token signing, and WebAuthn credentials, alongside stronger default SAML/WS-Federation assertion encryption.

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

01 **ML-DSA certificates for TLS, mTLS and RADIUS** NEW 78

authentik now generates and stores `ML-DSA-44`, `ML-DSA-65` and `ML-DSA-87` certificates alongside existing RSA, ECDSA, Ed25519 and Ed448 certificates, for use with TLS between components, mutual TLS, and RADIUS with EAP-TLS 1.3. Verifying these certificates requires outposts on 2026.11 or newer, and the feature is unavailable when FIPS compliance mode is on.

— Names cert types, use cases and FIPS caveat but no exact config path.

launch-20261103-8cbdf500

02 **ML-DSA signing for OAuth2 and OIDC tokens** NEW 73

OAuth2 providers can now sign ID tokens, access tokens and client assertions with an ML-DSA certificate; the proxy outpost and OAuth sources verify ML-DSA-signed tokens for federation with an upstream IdP that also signs with ML-DSA, and DPoP proofs can be bound to ML-DSA keys.

— Clear mechanism and scope, but no setting name given. launch-20261103-8cbdf500

03 **AES-256-GCM assertion encryption for SAML and WS-Fed** IMPROVED 72

SAML and WS-Federation providers now let you choose the assertion encryption algorithm, with new providers defaulting to `AES-256-GCM` while existing providers keep `AES-128-CBC` until changed; SAML sources now request `AES-256-GCM` from upstream IdPs.

— Names algorithms and defaults but not the exact config field. launch-20261103-8cbdf500

THINNER COVERAGE BELOW

04 **ML-DSA credentials in WebAuthn passkeys** NEW 50

WebAuthn registration and login now accept ML-DSA credentials, placed at the bottom of the accepted algorithm list so only an authenticator offering nothing else will use it.

— Describes behaviour but no actionable setting or command. launch-20261103-8cbdf500

WAS THIS USEFUL?  

## JumpCloud

Feedback ⓘ announced · Nov 2, 2026 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 launched Access Certification, a native identity governance capability for running scheduled or one-time user access review campaigns, complete with multi-step reviewer workflows, an admin task queue, tamper-evident audit reports, and Directory Insights automation hooks.

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

### 01 Directory Insights event triggers Workflows automation NEW

85

Directory Insights can emit an Access Certification event to trigger downstream automation via JumpCloud Workflows after a review concludes, and retains governance events for 90 days, so required audit evidence should be exported before that window expires.

Trigger follow-up automation — such as ticket creation or Slack notification — whenever a reviewer records a Revoke decision in a campaign.

📍 In JumpCloud Workflows, create a new workflow and set the trigger to the Access Certification event in Directory Insights, then add your follow-up action steps (e.g., call an external webhook or send a notification).

— Named event, retention limit, and exact workflow setup steps given.

launch-20261102-9f81ec04

### 02 Access Certification review campaigns NEW

75

JumpCloud adds Access Certification, a native identity governance and administration (IGA) capability for running scheduled or one-time access review campaigns inside the platform. Campaigns can target SSO applications or user groups, assign reviewers from administrators, a user's direct manager, resource owners, or specific user groups, and support multi-step reviews that route decisions through managers, resource owners, and reviewer groups at different stages of a single campaign. Remediation applies immediately after a final Revoke decision or on a schedule after the review phase closes, with no-response handling configurable to retain or revoke access; non-admin reviewers complete assignments in the JumpCloud User Portal, recording Keep or Revoke decisions individually or in bulk, with optional reassignment.

— Full mechanism and reviewer options named, but no API/flag to run directly.

launch-20261102-9f81ec04

THINNER COVERAGE BELOW

03 **Tamper-evident PDF audit reports** NEW 55

Export tamper-evident PDF reports documenting access review decisions, attestations, and administrator sign-off to support SOC 2, ISO 27001, and SOX audit preparation.

— Names output format and compliance frameworks but no export path.

launch-20261102-9f81ec04

04 **Admin Queue for pending review tasks** NEW 25

The JumpCloud Admin Queue consolidates pending review and remediation tasks assigned to an administrator.

— Bare description with no detail on scope or navigation. launch-20261102-9f81ec04

WAS THIS USEFUL?  

## RoboForm

Feedback ⓘ announced · Oct 30, 2026 BLOG ↗ RANK

RoboForm is a password manager and form-filling tool that securely stores login credentials and automatically fills web forms across browsers and devices.

> RoboForm 10.0.0 improves search, clipboard handling, and authentication on the desktop client.

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

01 **Auto-clearing clipboard after Editor copy** IMPROVED 55

The Editor now auto-clears the clipboard two minutes after copying data using the Editor's Copy button.

— Clear behaviour and timing specified launch-20261030-e004bc1e

02 **Folder results in Start Page search** IMPROVED 35

RoboForm's Start Page search now includes folders in its results, alongside existing logins and other items.

— Names the surface but no mechanism detail launch-20261030-e004bc1e

03 **Authentication and encryption updates** IMPROVED 20

RoboForm adds support for updated encryption and authentication settings, and updates the Desktop Unlock UI to use a new authentication protocol.

— No specifics on protocol or settings named launch-20261030-e004bc1e

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

# DNSFilter

Feedback ⓘ

11/03 November 3 – Web Application and API Update (2026-11-03) BLOG ↗ RANK

Give feedback to the DNSFilter team so we can make more informed product decisions. Powered by Canny.

> DNSFilter rolled out a new MSP branding system with custom CNAME subdomains and refined several roaming-client management UI flows, including inheritance labeling and a safer delete confirmation.

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

01 **Branding settings page replaces Whitelabel** NEW 69

DNSFilter adds a Branding settings page for MSPs supporting CNAME-based custom domains, replacing the legacy Whitelabel page. Found at `Settings > Branding` (previously `Settings > Whitelabel`), it lets MSPs upload logos and configure a managed branding subdomain, with existing assets carried over automatically.

When migrating a CNAME-based MSP white-label setup to the new managed branding subdomain, locate the replaced page in the dashboard.

> 📍 In the MSP dashboard, go to Settings › Branding (previously Settings › Whitelabel) to upload logos and configure your managed branding subdomain. Existing assets have been carried over automatically.

– Names exact UI path and migration behavior. 11/03 November 3 – Web Application and …

THINNER COVERAGE BELOW

02 **Roaming Client delete flow with uninstall and confirmation**

IMPROVED 52

The Roaming Client delete flow now distinguishes between Delete and Delete & Uninstall, with the latter option appearing only when uninstall-capable agents are selected. Single-client deletes now require type-to-confirm.

– Clear behavior change but no navigation detail given.

11/03 November 3 – Web Application and …

03 **Inherited policy/block page labeling in grids** IMPROVED 48

LAN Subnets, Relays, and Roaming Clients grids now show the inherited Policy or Block Page name with an Inherited from Site label, replacing the previous map-pin icon; global

policies are now marked with a globe icon beside the name.

— Describes exact UI change but no interaction steps.

11/03 November 3 – Web Application and …

WAS THIS USEFUL? 👍 👎

# amneziawg-installer

Feedback ⓘ v6.1.2 NOTES ↗ RANK

1 VPN server on a fresh Ubuntu or Debian VPS with one command. Made for networks that block WireGuard.

> amneziawg-installer adds terminal-based QR code generation for client onboarding and a standalone firewall guard against IP-options-based WireGuard detection.

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

## 01 IP-options firewall guard NEW

82

`sudo bash install_amneziawg.sh --ipopts-guard` deploys a standalone firewall rule that drops TCP packets with IP options, installing only the `raw` `-table` drop rule and the `amneziawg-ipopts-guard.service` without touching existing VPN settings. `manage check` reports whether this guard is active on a running server.

Add the IP-options firewall guard to an already-running server without reinstalling or reconfiguring the VPN.

```
$ sudo bash install_amneziawg.sh --ipopts-guard
```

— Names the exact flag, service, table and a status subcommand. v6.1.2

## 02 Terminal QR codes for client configs NEW

76

`manage qr <name>` prints a client config QR code directly in the terminal, eliminating the need to `scp` images off the server. The `--vpn` flag prints a `vpn://` deep-link QR for Amnezia VPN instead of the raw WireGuard config QR.

Print a client config QR code in the terminal to onboard a phone without copying files off the server.

```
$ manage qr my_phone
```

Print the `vpn://` deep-link QR for Amnezia VPN instead of the raw WireGuard config QR.

```
$ manage qr my_phone --vpn
```

— Exact commands and flag given; mechanism is simple but fully specified. v6.1.2

WAS THIS USEFUL? 👍 👎

## Firezone

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone launched an EU-region cloud deployment that keeps identity data, access policies, and audit logs stored exclusively within the European Union.

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

01 **EU-region cloud deployment** NEW 86

Firezone now offers an EU-region account option storing identity data, access policies, and audit logs exclusively in Sweden and France on Azure infrastructure, available on all plans including Starter. Clients switch to it via Settings → Control plane → Firezone Cloud (EU) before signing in, connecting to `wss://eu.api.firezone.dev` with the admin portal at `https://eu.app.firezone.dev`.

Connect a Firezone Client to the new EU region to keep all identity and policy data within the European Union.

> 📍 In the Firezone Client, go to Settings → Control plane → Firezone Cloud (EU), then sign in. The client will authenticate against wss:/ `/eu` . `api.firezone.dev` and the admin portal will be at `https://eu.app.firezone.dev`.

— Names exact endpoints, UI path, and data residency regions `launch-20261103-eb77fb80`

WAS THIS USEFUL? 👍 👎

## Twingate

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗ RANK

Twingate is a identity-based access for users, services, and AI agents that deploys in minutes, scales to every resource, and finally lets you retire your VPN.

> Twingate added an AI Usage Overview to its Secure DNS analytics, giving admins visibility into which AI tools employees access across DNS-filtered devices.

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

01 **AI Usage Overview in Secure DNS analytics** NEW 88

A new AI Usage Overview view inside Secure DNS analytics in the Admin Console, available to all DNS Filtering customers at no extra charge, shows which AI tools (including ChatGPT, Claude, Cursor, and Granola) employees access via DNS lookups. It includes an AI Usage by Date chart of daily AI service activity, an AI Exposure metric (AI requests as a share of all DNS requests or share of covered devices making AI requests), and a Top AI Services list rankable by Devices or Queries. Expanding a service row shows matched domains and per-domain request counts plus a Days Active column to distinguish one-off from recurring access, and device-level activity can be exported as a CSV for follow-up investigation.

Discover which AI services employees are accessing across DNS-filtered devices and check for unfamiliar tools.

> 📍 In the Twingate Admin Console, go to Secure DNS › AI Usage Overview. Select a time window and profile, then review the Top AI Services list. Switch between Devices and Queries views to rank services by reach or volume.

Investigate device-level activity for an unfamiliar AI service flagged in the overview by exporting a CSV for team follow-up.

> 📍 In the Twingate Admin Console, go to Secure DNS › AI Usage Overview. Expand the row for the service under investigation, then export device-level activity as a CSV.

— Names exact UI metrics, columns and CSV export with clear navigation path.

launch-20261103-d24acdbc

WAS THIS USEFUL? 👍 👎

## EVALUATE

◆ AI Observability & Evals

# ongrid

Feedback ⓘ v0.20.0 NOTES ↗ RANK

An ops AI Agent that understands your infrastructure, finds the root cause, and fixes it — right from Slack, Telegram, Lark or DingTalk.

> ongrid's v0.20.0 release bundles new agent tooling, IM approvals, SNMP support and edge upgrade automation alongside a floating AIOps assistant with page context awareness and Elasticsearch Basic Auth support for log pipelines.

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

01 **Elasticsearch Basic Auth log destination** NEW 43

ongrid now supports Elasticsearch with Basic Auth as a log destination for its log pipelines.

— Names auth method and destination type, no config key given v0.20.0

02 **Floating AIOps assistant with page context** NEW 30

ongrid adds a floating AIOps assistant that is aware of the page context it is displayed on.

— Describes behavior but no UI path or config detail v0.20.0

03 **Agent tooling, IM approvals, SNMP support and edge upgrades** NEW 22

ongrid adds new agent tooling, approval workflows over instant messaging, SNMP protocol support, and automation for edge upgrades.

— Four capabilities named with no mechanism or surface detail v0.20.0

WAS THIS USEFUL? 👍 👎

## GOVERN

◆ AI/LLM Security

# NVIDIA OpenShell

Feedback ⓘ v0.1.7 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell's v0.1.7 release focuses on telemetry and networking hardening — adding OTLP trace forwarding from sandboxes, mutual-TLS and private-PKI options for gateway traffic, direct-to-replica SDK routing, and a switch to OCSF 1.8.0-only logging that deprecates older schema settings.

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

## 01 OCSF 1.8.0 becomes the only emitted schema

BREAKING

85

OpenShell now emits only OCSF 1.8.0 records; deployments that relied on 1.1 or 1.3 records will now receive 1.8.0 instead. The `ocsf_schema_version` setting, the gateway `schema_version` fields, Helm's `server.ocsfLog.schemaVersion`, and `OPENSHELL_OTLP_OCSF_SCHEMA_VERSION` are deprecated and ignored, and setting any of them logs a warning.

— Lists every deprecated setting and the resulting behaviour clearly v0.1.7

## 02 Expanded OTLP transport and auth options

IMPROVED 80

OTLP export can now use HTTP instead of gRPC by setting `protocol = "http/protobuf"` on `[openshell.gateway.otlp]` or any signal sub-table (gRPC remains default); requests can be authenticated or routed with `headers` or `headers_file` on the same tables, letting collectors like MLflow receive traces via `x-mlflow-experiment-id`; and collectors can be reached over mutual TLS or private CAs by setting `ca_file`, and optionally `cert_file` and `key_file`.

Connect your OTLP collector over mutual TLS using a private CA, so traces are encrypted end-to-end without a public certificate.

toml

```toml
[openshell.gateway.otlp]
ca_file = "/etc/openshell/pki/ca.pem"
cert_file = "/etc/openshell/pki/client.pem"
key_file = "/etc/openshell/pki/client-key.pem"
```

— Names every OTLP config key with a worked mTLS example v0.1.7

## 03 Sandbox agent trace forwarding to OTLP collectors

NEW

75

Sandboxed agent OpenTelemetry traces can now be forwarded to an external collector without writing an egress rule, by adding `[openshell.gateway.otlp.agent]` with `enabled = true` to `gateway.toml`; the supervisor tags each span with the sandbox id before forwarding.

Send agent traces from inside a sandbox to your OTLP collector, authenticated with an MLflow experiment header, without writing an egress rule.

```toml
[openshell.gateway.otlp.agent]
enabled = true

[openshell.gateway.otlp]
protocol = "http/protobuf"
headers_file = "/etc/openshell/otlp-headers.txt"
```

— Config key plus mechanism (span tagging) and a runnable example v0.1.7

04 **Containerized CLI browser login fix** NEW 75

Running the CLI inside a container with `OPENSHELL_CALLBACK_PORT` and `OPENSHELL_CALLBACK_BIND_ADDRESS` set now fixes the OAuth callback so browser login completes correctly, e.g. `docker run --rm -e OPENSHELL_CALLBACK_PORT=9876 -e OPENSHELL_CALLBACK_BIND_ADDRESS=0.0.0.0 -p 9876:9876 my-openshell-cli openshell gateway login`.

Run the CLI inside a container and fix the OAuth callback so browser login completes correctly.

```
$ docker run --rm -e OPENSHELL_CALLBACK_PORT=9876 -e
OPENSHELL_CALLBACK_BIND_ADDRESS=0.0.0.0 -p 9876:9876 my-
openshell-cli openshell gateway login
```

— Exact env vars and a complete runnable docker command v0.1.7

05 **Private PKI for sandbox egress** NEW 70

Sandbox egress can now be routed through a private PKI by setting `additional_ca_bundle` under `[openshell.supervisor.network]` to a PEM bundle; this applies to every first-party compute driver.

— Exact config key and scope given, but no worked example v0.1.7

06 **Direct-to-replica exec and tunnel routing in SDKs** IMPROVED

70

The Rust, Go, Python, and TypeScript SDKs now send exec calls directly to the gateway replica that owns the sandbox when replica routing is enabled; the Go SDK also routes port-forward and SSH tunnel calls this way, and the TypeScript SDK routes port-forward calls.

— Names all four SDKs and which call types each routes v0.1.7

---

07 **Preflight validation for supervisor pool config** BREAKING

60

`config preflight` and gateway startup now reject an `[openshell.gateway.supervisor_pool]` table when the selected driver does not use shared supervision, and also reject a pool recovery timeout shorter than member failure detection.

— Names the command and config table but no remediation example v0.1.7

THINNER COVERAGE BELOW

08 **Hashed hostnames for shared-cluster gateway services** NEW

55

When multiple gateways share a cluster, each sandbox service can now get a unique hostname by setting `service_hostname_format = "hashed"` under `[openshell.gateway]`; the previous `readable` default is unchanged.

— Names the exact config key and both values v0.1.7

---

09 **Direct supervisor TLS listener for sandbox services** NEW

40

A new supervisor TLS listener lets declared HTTP and WebSocket sandbox services accept application traffic directly, bypassing the gateway. It is disabled by default.

— No config key named, only prose description of the feature v0.1.7

---

10 **Typed Python gRPC stubs** IMPROVED

40

The Python gRPC bindings now ship typed sync and async stubs, giving editors and type checkers visibility into concrete request and response types.

— Describes the benefit but no names of specific stub types v0.1.7

└──▷ **BREAKING ON UPGRADE**

! OpenShell now emits only OCSF 1.8.0 records; deployments that relied on 1.1 or 1.3 records will now receive 1.8.0 records. The `ocsf_schema_version` setting, the gateway `schema_version` fields, Helm `server.ocsfLog.schemaVersion`, and `OPENSHELL_OTLP_OCSF_SCHEMA_VERSION` are deprecated and ignored; setting them logs a warning.

! `config preflight` and gateway startup now reject an `[openshell.gateway.supervisor_pool]` table when the selected driver does not use shared supervision, and also reject a pool recovery timeout shorter than member failure detection.

---

WAS THIS USEFUL? 👍 👎

# Docker Sandboxes

Feedback ⓘ v0.51.0 NOTES ↗ RANK

Docker Sandboxes run AI coding agents in isolated environments on your machine or on Docker-managed cloud infrastructure

> Docker Sandboxes v0.51.0 adds custom workspace mounts, FUSE filesystem support, stdin secret input, and new policy and MCP controls, alongside a breaking migration of built-in agent kits to published v3 registry images and several deprecated template images.

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

## 01 Built-in kits move to published v3 registry images

BREAKING

90

Built-in agent names now resolve to published kits at `docker.io/sbx/<name>:latest` and the `kit.allowExtractedAgents` setting is removed. Sandboxes created from embedded Claude or Codex specifications must be recreated before using `sbx kit add`, `sbx run --model`, or `sbx move` — save files and unpushed commits before removing. The shell agent now uses a published Debian 13 v3 kit and requires registry access on first use, so add the `shell-credentials` mixin to provide provider credentials. V1 and v2 kits cannot be combined with v3 kits; older mixins must be upgraded to v3 before combining with the new built-in kits.

— Full migration path with exact commands and settings named v0.51.0

## 02 Read secrets from stdin NEW

80

Read secrets from standard input with `sbx secret set-custom --value-stdin`, keeping secret values out of command arguments and shell history.

Keep secrets out of shell history by piping a token directly into a custom secret.

```
$ echo 'my-token-value' | sbx secret set-custom --value-stdin
```

— Runnable example shows exact command v0.51.0

## 03 New settings to disable telemetry and bundled model server

NEW

80

Disable local CLI and daemon telemetry persistently with `sbx settings set telemetry.disabled true`, or disable the bundled model server and its inference-runtime download by setting `model.server.enabled` to `false`; run `sbx daemon restart` to apply either change to a running daemon.

Disable telemetry collection on a CI machine so neither the CLI nor the daemon reports usage.

```
$ sbx settings set telemetry.disabled true && sbx daemon restart
```

— Runnable commands cover full disable workflow Names exact setting and restart step

v0.51.0

04 **Custom workspace mount targets** NEW 75

Mount additional workspaces at custom paths by passing `PATH:TARGET[:ro]` to `sbx create` and `sbx run`, or by setting `additionalWorkspaces[].target` in environment files.

— Names exact flag syntax and config key v0.51.0

05 **MCP tool naming and static mode controls** NEW 70

Prefix MCP tool names with their server name by enabling `mcp.useServerNamePrefixes`, and run sandboxes in static MCP mode without attaching servers by passing `--static-mcp none` to `sbx create` or `sbx run`.

— Names both the config key and the flag v0.51.0

06 **Sandbox template images deprecated for published kits**

DEPRECATED 70

The `docker/sandbox-templates` images for `codex`, `cursor-agent`, `devin`, `docker-agent`, and `opencode` (including `-docker` variants) are no longer published. The `docker/sandbox-templates:shell` and `:shell-docker` images are no longer refreshed; use `docker.io/sbx/shell-image` and `docker.io/sbx/shell-docker-image` for custom workloads.

— Names all deprecated images and their replacements v0.51.0

07 **Start agents in a subdirectory via --workdir** NEW 65

Start agents in a subdirectory with `sbx run --workdir`, including when reattaching to an existing sandbox.

— Clear flag named, little mechanism detail v0.51.0

08 **Suppress creation progress output** NEW 65

Suppress creation progress output while still receiving command results using `sbx run --no-progress`, or `sbx create --no-progress` for creation.

— Names both flags, behavior is simple v0.51.0

09 **Custom secrets now require explicit destination hosts**

65

BREAKING

Custom secrets without destination hosts no longer substitute credentials. Existing mappings must be repaired by adding explicit destination hosts while preserving their placeholders and other settings.

— Gives migration step but no exact field names v0.51.0

10 **FUSE filesystem support in local sandboxes** NEW 60

Enable FUSE filesystems in local sandboxes via the `sandbox.fuse` setting or the environment option `sandboxOptions.fuse`.

— Names config keys but not mechanism v0.51.0

THINNER COVERAGE BELOW

11 **HTTP governance policy check command** NEW 55

Check HTTP governance policy behavior with the new `sbx policy check http` command.

— Bare command name, no behavior detail v0.51.0

12 **Clipboard copy requires opt-in setting** BREAKING 55

Copying text from a local sandbox to the host clipboard now requires opting in with the `clipboard.textCopy` setting.

— Names setting but no further mechanism given v0.51.0

└──▷ **BREAKING ON UPGRADE**

! Built-in agent names now resolve to published kits at `docker.io/sbx/<name>:latest`; the `kit.allowExtractedAgents` setting is removed.

! Sandboxes created from embedded Claude or Codex specifications must be recreated before using `sbx kit add`, `sbx run --model`, or `sbx move` — save files and unpushed commits before removing.

! The shell agent now uses a published Debian 13 v3 kit and requires registry access on first use; add the `shell-credentials` mixin to provide provider credentials.

! V1 and v2 kits cannot be combined with v3 kits — older mixins must be upgraded to v3 before combining with the new built-in kits.

! The `docker/sandbox-templates` images for `codex`, `cursor-agent`, `devin`, `docker-agent`, and `opencode` (including `-docker` variants) are no longer published.

! The `docker/sandbox-templates:shell` and `:shell-docker` images are no longer refreshed; use `docker.io/sbx/shell-image` and `docker.io/sbx/shell-docker-image` for custom workloads.

! Copying text from a local sandbox to the host clipboard now requires opting in with the `clipboard.textCopy` setting.

! Custom secrets without destination hosts no longer substitute credentials — existing mappings must be repaired by adding explicit destination hosts while preserving their placeholders and other settings.

---

WAS THIS USEFUL? 👍 👎

# h5i

Feedback ⓘ v0.5.3 NOTES ↗ RANK

Build, hack, and verify web apps. Find vulnerabilities through red-teaming.

> h5i's v0.5.3 release adds two new workbench plugins for web app red-teaming (websec and recon), verified TLS support for PostgreSQL connections, and expanded Lean 4 formal verification coverage for integer columns.

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

01 **Websec workbench plugin for HTTP traffic manipulation** NEW 75

Install via `h5i plugin install websec --from ./h5i-websec` to capture, modify, replay, and compare HTTP requests through the same policy and budget as the browser.

— Names exact install command and policy/budget mechanism v0.5.3

02 **Recon plugin for endpoint discovery** NEW 68

Install via `h5i plugin install recon --from ./h5i-recon` to record endpoints, parameters, and candidate URLs a target exposes during discovery.

Install the recon plugin after downloading its archive to start discovering endpoints and hidden attack surfaces on a target.

```
$ curl -L https://github.com/h5i-dev/h5i/releases/download/v0.5.3/h5i-recon-v0.5.3-x86_64-unknown-linux-musl.tar.gz | tar -xz
h5i plugin install recon --from ./h5i-recon
```

— Exact install command plus worked download-and-install example v0.5.3

THINNER COVERAGE BELOW

03 **Verified TLS to PostgreSQL in h5i-app** NEW 45

h5i-app now supports verified TLS to PostgreSQL via the `tls` feature flag.

— Names the flag but no further mechanism detail v0.5.3

04 **Lean 4 lemmas for u32 columns** NEW 35

h5i-app-schema now includes Lean 4 lemmas for u32 columns to extend formal verification coverage, building on existing work for range comparisons on integer columns in h5i-app reads.

— Names the module and data type but no usage path v0.5.3

### └──▷ ALSO FROM THESE RELEASES

Replay a suite of attack flows against a staging target and fail CI if OpenAPI coverage drops below 80%.

```
$ h5i test .h5i-tests/tests --target https://staging.example.com --openapi openapi.yaml --min-coverage 80
```

---

WAS THIS USEFUL? 👍 👎

# ToolHive

Feedback ⓘ v0.60.0 NOTES ↗ RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive's v0.60.0 release tightens LLM authentication controls and extends its token exchange subsystem with Redis-backed storage, refresh token support, and stricter trust requirements for dynamic client registration.

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

01 **Scoped LLM authentication via --scopes flag** NEW 83

The `thv llm setup` and `thv config set` commands now accept a `--scopes` flag to restrict which OAuth scopes are requested when registering an LLM provider, for example `thv llm setup --scopes openid,profile,email`.

Restrict which OAuth scopes are requested when registering an LLM provider, so tokens carry only the permissions your deployment needs.

```
$ thv llm setup --scopes openid,profile,email
```

— Exact flag and runnable example command provided v0.60.0

THINNER COVERAGE BELOW

02 **Required trust mode for private_key_jwt delegates** BREAKING 50

A trust mode is now required for `private_key_jwt` delegates, enforcing explicit trust configuration before asymmetric client authentication is permitted.

— Names auth method but no exact config key for trust mode v0.60.0

03 **Redis storage backend for token exchange grants** NEW 45

OAuth 2.0 exchange grants can now be stored in Redis by configuring the new Redis storage backend for the token exchange subsystem.

— Names Redis backend but no config key or flag given v0.60.0

04 **Configurable token endpoint auth for DCR upstreams** IMPROVED 40

Dynamic client registration (DCR) upstreams can now choose their token endpoint authentication method, expanding the range of identity providers that work with

dynamic client registration.

— Describes behaviour but no concrete config surface named v0.60.0

05 **Refresh token issuance through token exchange** NEW

35

Token exchange can now issue refresh tokens, letting delegate clients obtain long-lived credentials without needing to re-authenticate.

— No flag or config key named for this capability v0.60.0

WAS THIS USEFUL? 👍 👎

## PLATFORMS & UTILITIES

◆ General Utilities & Libraries

## OpenSSL

Feedback ⓘ announced · Oct 23, 2026 BLOG ↗ RANK

General purpose TLS and crypto library

> OpenSSL 4.1 lands DTLS 1.3 support alongside GREASE, IKEV2 KDF, and SSL listener API updates, plus a round of ML-DSA/ML-KEM and AES-CBC performance optimizations, new architecture and build-target support, and several breaking changes to Configure options and legacy build targets.

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

01 **DTLS 1.3 support (RFC 9147)** NEW 60

OpenSSL now supports DTLS 1.3 per RFC 9147, documented in the new ossl-guide-dtlsv13(7) manual page.

– Names RFC and manual page but no usage example launch-20261023-bf9abaaa

02 **Cryptographic performance optimizations across architectures** IMPROVED 60

OpenSSL adds optimized ML-DSA and ML-KEM NTT operations on `ppc64le`, optimized ML-DSA operations on `x86_64`, AVX-512-optimized SHAKE x4 for ML-DSA on `x86_64`, and AVX-512 and VAES optimizations for AES-CBC decryption on `x86_64`.

– Names several optimizations and architectures, no benchmarks given launch-20261023-bf9abaaa

03 **tsget migrates to Net::Curl::Easy** BREAKING 60

The `tsget` utility now uses `Net::Curl::Easy` (from the Net-Curl CPAN distribution) instead of the abandoned `WWW::Curl::Easy`; users must install `Net::Curl::Easy` before upgrading.

– Names exact packages and gives clear pre-upgrade action launch-20261023-bf9abaaa

04 **no-ecdsa and no-ecdh removed from Configure** BREAKING 60

The `no-ecdsa` and `no-ecdh` options are removed from Configure; use `no-ec` to disable elliptic curve support.

– Names removed flags and the exact replacement flag launch-20261023-bf9abaaa

THINNER COVERAGE BELOW

05 **MSVC 2013 build targets** NEW 50

OpenSSL adds `VC-WIN32-MSVC2013` and `VC-WIN64A-MSVC2013` build targets to bridge C99 standard gaps in MSVC 2013.

— Exact target names given, usable directly in build config launch-20261023-bf9abaaa

06 **Legacy Windows build targets removed** BREAKING 45

The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are removed from Configurations.

— Names removed targets but no migration path given launch-20261023-bf9abaaa

07 **RFC 8701 GREASE support in TLS** NEW 35

OpenSSL now supports RFC 8701 GREASE (Generate Random Extensions And Sustain Extensibility) in TLS.

— Names the RFC but no mechanism or config detail launch-20261023-bf9abaaa

08 **DTLS support in SSL listener API** NEW 35

OpenSSL now includes DTLS support in the SSL listener API.

— Bare statement with no API detail or example launch-20261023-bf9abaaa

09 **Initial Elbrus2000 (e2k) architecture support** NEW 25

OpenSSL now adds initial support for the Elbrus2000 ( `e2k` ) architecture.

— Names architecture but no build or usage detail launch-20261023-bf9abaaa

10 **IKEV2 KDF support** NEW 20

OpenSSL now supports the IKEV2 KDF.

— Single-line mention with no further specifics launch-20261023-bf9abaaa

### └──▷ BREAKING ON UPGRADE

! The `tsget` utility now uses `Net::Curl::Easy` (from the Net-Curl CPAN distribution) instead of the abandoned `WWW::Curl::Easy` ; users must install `Net::Curl::Easy` before upgrading.

! The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are removed from Configurations.

! The `no-ecdsa` and `no-ecdh` options are removed from `Configure` ; use `no-ec` to disable elliptic curve support.

WAS THIS USEFUL?  

## INFRASTRUCTURE

◆ Internet Services

# Cloudflare

Feedback ⓘ WAF - WAF Release - 2026-11-03 (2026-11-03)

BLOG ↗ RANK

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

> Cloudflare expanded its WAF Managed Ruleset with two new detections: a phpMyAdmin authentication bypass rule and a scheduled Dangerous File Upload rule rolling out in Log mode.

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

01 **New detections added to WAF Managed Ruleset** NEW 47

Cloudflare WAF adds a new detection for phpMyAdmin authentication bypass in the Cloudflare Managed Ruleset (rule `...13f91633`), and a new Generic Rules - Dangerous File Upload detection (rule ID `...7178209b`), the latter deploying in Log mode on 2026-11-10.

— Names exact rule IDs and deployment mode but no further mechanism.

WAF - WAF Release - Scheduled changes f… WAF - WAF Release - 2026-11-03 (2026-11…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 4, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue