

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 21, 2026 · EVERY DAY

EDITIONS **tail** | grep | head | ls | diff | uniq

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 163 changes

Do you prefer **scroll**?



▸ // CONTENTS SHOW

50 tools

▼ // THE BRIEF HIDE

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

PROTECT

Block a malicious dependency before the runner installs it

vim ? help

GitHub's preview firewall sits in front of npm fetches for CI/CD runners and coding agents, which is the window where a typosquat or compromised release actually lands on a build host; JFrog adds a semantic scanner and Agent Guard for malicious agent skills and plugins, a surface SCA and SAST have no signatures for. Socket's async Fixes API endpoints and authorization checks let remediation run as a job rather than a blocking call.

[GitHub](#) · [JFrog](#) · [Socket](#)

DETECT

Find the account that was already taken over before you onboarded

Managed ITDR Lookback scans 180 days of Microsoft 365 sign-in history automatically at tenant onboarding, closing the usual blind spot where monitoring starts at day zero and a months-old mailbox rule or persistent session is invisible. The on-demand Identity Posture Assessment ships as a shareable PDF, which is what a client conversation actually needs.

[Huntress](#)

PROTECT

Catch the auth bug no SAST rule pattern-matches

The early-access Business Logic Security Analyzer targets broken authorization and mass-assignment, flaws that are invisible to taint tracking because the code is syntactically correct and only wrong against intent. These are the findings that normally require a human reviewer or a pentest cycle.

[GitLab](#)

GOVERN

Account for which AI agents are running and who approved them

Nudge extends agent discovery to agents built on Notion and ties findings into ServiceNow and Jira Cloud with auto-close on resolution, so the remediation record lives where the service desk already works instead of in a separate console. JumpCloud adds REST APIs for AI Agent management, which is the prerequisite for treating agents as managed identities rather than shadow accounts.

[Nudge Security](#)

◆ **Reconnaissance & OSINT**

nerva

[Feedback](#)[v1.74.0](#)[NOTES ↗](#)[STABLE](#)

Fast service fingerprinting CLI for 170+ protocols (TCP/UDP/SCTP) - built by Praetorian

↳ [GET THIS VERSION](#)

```
$ git clone --branch v1.74.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.74.0
```

[NEW DETECTION](#) [CUPS fingerprinter](#)

[NEW DETECTION](#) [SHOUTcast DNAS fingerprinter](#)

▼ **1 more release in this issue** · 2026-10-20

[v1.73.0](#)[NOTES ↗](#)[STABLE](#)

↳ [GET THIS VERSION](#)

```
$ git clone --branch v1.73.0 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.73.0
```

WAS THIS USEFUL?



ESP32 Marauder

Feedback



v1.19.0

NOTES ↗

RANK

A suite of WiFi/Bluetooth offensive and defensive tools for the ESP32.

ESP32 Marauder v1.19.0 adds support for the Seeed XIAO ESP32-C5 board, upgrades wardrive logging to the WiGLE 1.6 format, and adds battery fuel gauge IC support for existing hardware revisions.

↳ WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

01 Seeed XIAO ESP32-C5 hardware support NEW 57

Adds the Seeed XIAO ESP32-C5 as a supported hardware target, shipping a dedicated firmware binary `_xiao_esp32c5.bin`, and adds corresponding C5 hardware aliases to the installer catalog.

– Names the binary and target but not install steps v1.19.0

02 Battery fuel gauge IC support NEW 50

Adds support for the `IP5306` and `MAX17048` battery fuel gauge ICs on V4/V6/V6.1 hardware.

– Names ICs and hardware versions but no behaviour details v1.19.0

03 WiGLE 1.6 wardrive log format IMPROVED 25

Upgrades wardrive logs to the WiGLE 1.6 format.

– Bare mention of format version, no further detail v1.19.0

WAS THIS USEFUL?



Rapid7 Metasploit Pro

Feedback



October 19, 2026: 5.1.0-2026101901

NOTES ↗

RANK

Metasploit Pro is a commercial penetration testing platform from Rapid7 that automates exploitation, post-exploitation, and vulnerability validation for security assessments.

Metasploit Pro 5.1.0 adds 20 new modules this cycle, headlined by exploit modules for AI/LLM platforms (LiteLLM, OpenCTI) and other CVEs, a set of MongoDB enumeration and credential-harvesting modules, Windows AArch64 fetch adapters, and new persistence/post-exploitation modules touching Ollama and Linux PAM.

↳ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 AV/EDR detection in enum_protections

IMPROVED

75

Updates `post/linux/gather/enum_protections` to detect AV/EDR software present on the target system.

Check a compromised Linux host for active AV/EDR products after gaining a session.

```
$ use post/linux/gather/enum_protections
set SESSION 1
run
```

– Named module with a runnable usage example

October 19, 2026: 5.1.0-2026101901

02 New exploit modules

NEW

75

Adds exploit modules for CVE-2026-42271 (command execution in BerriAI LiteLLM proxy MCP test REST endpoints), a chain of CVE-2026-27960 and CVE-2026-39980 for unauthenticated remote command execution in OpenCTI, a SafeJS sandbox escape achieving RCE as root inside the OpenCTI API container (affecting versions 6.8.16 and above), CVE-2026-53959 (broken access control in 4gaBoards versions 3.3.8 and below exposing the `/api/users` REST endpoint and fields including `sssoGoogleEmail` and `sssoGithubEmail`), EDB-52079 (unauthenticated RCE in dizqueTV, a Node.js IPTV server), and CVE-2024-9264 (authenticated DuckDB command injection in Grafana SQL Expressions).

– CVE and endpoint named but no module path or command Describes CVE chain and impact but no module name Mechanism, privilege level and version scope given, no command Named CVE, version range, endpoint and exposed fields Names advisory and product but no further mechanism CVE and feature area named, no command or steps given

October 19, 2026: 5.1.0-2026101901

03 MongoDB enumeration and credential modules NEW 68

Adds `mongodb_version`, `mongodb_login`, `mongodb_schemadump`, and `mongodb_hashdump` auxiliary modules for MongoDB enumeration and credential harvesting.

– Names four modules but gives no example invocation October 19, 2026: 5.1.0-2026101901

04 PAM shared object post-exploitation module NEW 60

Adds a post module that uploads a PAM shared object file into the Linux authentication chain, compiling it on-target when the system is not x64.

– Clear mechanism described but no module name or command given

October 19, 2026: 5.1.0-2026101901

THINNER COVERAGE BELOW

05 Ollama auto-update persistence module NEW 45

Adds a Windows persistence module that leverages the Ollama auto-update mechanism.

– Explains mechanism but no module name or command October 19, 2026: 5.1.0-2026101901

06 Windows AArch64 fetch adapters NEW 40

Adds Windows AArch64 fetch adapters, extending payload delivery to ARM64 Windows targets.

– States the addition but no mechanism or usage detail October 19, 2026: 5.1.0-2026101901

WAS THIS USEFUL?  

APPLICATION SECURITY

Socket

Feedback



changelog · Oct 21, 2026

CHANGELOG

RANK

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

Socket shipped new async endpoints and authorization checks for its Fixes API, added token identification fields to the token-list API, and moved organization settings into the dashboard sidebar — alongside a breaking change to the format of masked API keys.

← WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Token identification fields in token-list API NEW

65

Token-list API responses now include `token_prefix` and `token_last4` fields, allowing key identification without parsing the masked display string.

– Names exact new fields but not full endpoint path `changelog-20261021-d33541bf`

02 Repository access checks for fix endpoints NEW

60

Adds repository access checks before computing or returning fixes for user OAuth sessions governed by repository or workspace policies, returning `403` for unfiltered repository endpoints.

– Names status code and policy scope but not endpoint path `changelog-20261021-d33541bf`

03 New masked format for API keys

BREAKING

60

API keys now use a new consistent masked format (e.g. `sktsec_...Xy9Q`); integrations that parse the previous masked format may need updating.

– Gives example format and migration implication `changelog-20261021-d33541bf`

THINNER COVERAGE BELOW

04 Async fix computation endpoints in Fixes API NEW

55

The Fixes API now documents async endpoints for starting a fix computation and polling for results, supporting `supported parameters` and result format to avoid timeouts on large projects or first runs.

– Describes mechanism but no exact endpoint paths given `changelog-20261021-d33541bf`

05 Organization settings link in dashboard sidebar NEW

Organization settings are now accessible directly from the [Settings](#) link at the bottom of the dashboard sidebar.

– Clear UI navigation path given `changelog-20261021-d33541bf`

↳ BREAKING ON UPGRADE

! API keys now use a new consistent masked format (e.g. `sktsec_...Xy9Q`); integrations that parse the previous masked format may need updating.

WAS THIS USEFUL?



RANK

JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

JFrog's biggest moves this window are on AI supply-chain security: Shadow AI detection now reaches developer workstations via a new Agent Plugin, and a new semantic scanner plus Agent Guard enforcement component catches malicious agent skills and plugins that SCA/SAST miss. JFrog Curation also gained three new governance capabilities — Compliant Version Selection, Curation Federation, and On-Demand Curation.

←→ WHAT SHIPPED • 5 FEATURES most completely described first ? what's the number?

01 Semantic AI asset scanning and enforcement with Agent Guard

NEW

84

JFrog AI Security performs semantic (intent-based) analysis of agent skill files (e.g. **SKILL.md**), agent plugins, and APM manifest packages to detect prompt injection, hidden instructions, obfuscated payloads, and credential-exfiltration instructions that SCA and SAST cannot reach. **JFrog Agent Guard** enforces deterministic-plus-semantic policy at the agent harness — the point an agent actually invokes the asset — not only at download time, so a skill reclassified after initial approval is blocked on machines that already have it. Each asset receives one of three classifications — **Not Malicious**, **Suspicious**, or **Malicious** — routed to block, review queue, or allow; per-version allow-lists mean approving one version does not auto-approve later versions; enforcement supports severity thresholds (e.g. block any skill with a Critical-severity finding) and time-bound waivers with fixed expiration dates. Agent skills, plugins, and APM packages are stored in **JFrog Artifactory** alongside Docker images and npm packages under the same permissions model and audit trail.

— Names components, classifications, and controls but no exact API or command

Scanning AI Assets: Catching What SCA a...

02 Shadow AI detection on developer workstations NEW

78

The **JFrog Agent Plugin**, installable on developer workstations, discovers and locally scans AI assets — skills, MCP servers, plugins, hooks, commands, and monitors — that bypass central repository controls. It runs a three-step pipeline: Discover (enumerate AI assets), Scan Locally (deterministic rules-based scan), and Report (send metadata,

findings, and risk analysis to the JFrog Platform). Results populate a unified AI asset inventory in JFrog AI Security showing each asset's name, type, reporting harness (Claude Code, Cursor, etc.), workstation count, local risk analysis, JFrog origin, and detection timestamp, with risk classifications for Shell Access + WebFetch (high exfiltration risk), Unscoped Write/Edit (writes outside designated directories), and BypassPermissions Mode (executes without required approval), plus an Origin field showing whether the asset matches one already managed in JFrog Artifactory or JFrog Catalog, or has **No Match**.

– Describes full mechanism and UI fields but no exact install command

Shadow AI Detection Now Extends to Deve...

THINNER COVERAGE BELOW

03 **On-Demand Curation for uncatalogued repositories** **NEW**

55

On-Demand Curation extends JFrog Curation policy enforcement to packages from internal, private, or ecosystem-specific repositories not covered by the JFrog Public Catalog, applying real-time vulnerability, license, immaturity, and malicious code evaluation before a package reaches the developer.

– Explains scope and checks performed but no config or UI path given [product docs](#)

04 **Curation Federation across Platform Deployment sites** **NEW**

45

Curation Federation synchronizes policies, conditions, and JFrog Catalog labels from a central controller site to all connected JFrog Platform Deployment sites, eliminating manual policy replication.

– Names the sync mechanism and scope but no setup steps given [product docs](#)

05 **Compliant Version Selection in JFrog Curation** **NEW**

42

When a requested package version is blocked by curation policy, JFrog Curation automatically returns the highest package version that passes all policies, letting developers avoid failed builds without manual intervention.

– Clear mechanism but no concrete flag, config, or endpoint named [product docs](#)

WAS THIS USEFUL?  

Endor Labs

[Feedback](#)

Your model, our analysis: BYOK for Endor Labs AI SAST | Endor Labs (2026-10-20)

BLOG »

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

Endor Labs AI SAST adds BYOK support routing model requests through customer-owned LLM provider accounts, with an AI Gateway enforcing encrypted key store, provider allowlist, and tenant isolation.

NEW SUPPORTED PLATFORM

Adds **BYOK** (bring your own key) support for AI SAST and AI Chat in the Endor Labs AURI platform, routing all model requests through your own OpenAI, Azure OpenAI, Gemini, or Vertex AI account instead of Endor Labs' credentials — with no fallback to Endor Labs if a key is configured.

Supports passing provider API keys as environment variables during AI SAST scans so keys are never stored with Endor Labs.

Adds an Endor Labs AI Gateway that enforces an encrypted key store, a provider allowlist, and tenant isolation for all BYOK model requests.

Adds custom organizational context ingestion for AI SAST — including Infrastructure as Code files, repository files such as **AGENTS.md** and **CLAUDE.md**, triage feedback, and uploaded threat models — to inform analysis with deployment and triage history.

NEW SUPPORTED PLATFORM

C language support for AI SAST

WAS THIS USEFUL?



Sigstore cosign

[Feedback](#)[v3.1.4](#)[NOTES ↗](#)[RANK](#)

Code signing and transparency for containers and binaries

cosign v3.1.4 adds a new subcommand for creating Sigstore bundles directly from container images, plus a minor shell completion improvement.

↳ [WHAT SHIPPED](#) · [2 FEATURES](#) most completely described first [? what's the number?](#)

01 Create bundle from container image NEW 75

The new `cosign bundle create-from-container` subcommand creates a Sigstore bundle directly from a container image, enabling offline or air-gapped verification workflows.

Create a Sigstore bundle from an already-signed container image for offline or air-gapped verification workflows.

```
$ cosign bundle create-from-container $IMAGE
```

– Named subcommand with example command, but mechanism detail thin v3.1.4

THINNER COVERAGE BELOW

02 Shell completion for --slot option IMPROVED 35

Adds shell completion for the `--slot` option argument.

– Named flag but no detail on behaviour or usage v3.1.4

WAS THIS USEFUL?



Openlayer

[Feedback](#)


Governance lifecycle stages, bias and fairness tests, and multiple project dashboards

[NOTES](#)
[RANK](#)

Openlayer is the AI governance and observability platform for regulated enterprises: one place to test, monitor, and govern every model, agent, and AI system your company runs.

Openlayer documented its gateway's full error-handling and request-tracing behavior while shipping a wave of governance and testing capabilities — bias/fairness, prompt injection, and data freshness tests, per-step evaluation, custom JSONPath columns, expanded CLI tooling, API key lifecycle management, and new governance registry endpoints.

←→ **WHAT SHIPPED · 8 FEATURES** most completely described first [? what's the number?](#)

01 Gateway error handling and request tracing documented NEW

90

Gateway requests can now carry an **X-Request-ID** header, which is echoed back as the Openlayer trace ID for log correlation. Error formats follow the upstream provider shape per endpoint: `/v1/messages` returns the Anthropic error shape, `/v1/responses` and `/v1/embeddings` return the OpenAI shape, and `/v1/models` returns a shape based on path (`/anthropic/v1/models` or `/openai/v1/models`) or SDK headers when called at the bare `/v1/models` . HTTP status codes map to error **type** values: **invalid_request_error** (400/422), **authentication_error** (401), **permission_error** (403), **not_found_error** (404), **rate_limit_error** (429), **api_error** (5xx). 503 responses carry **Retry-After** headers that vary by cause: messages prefixed **gateway at capacity** and **no upstream capacity for** carry 1s, guardrail-loading and database-unavailable errors carry 5s, and provider config errors carry none. `/v1/images/generations` errors return `{"detail": ...}` with guardrail blocks including **type** , **stage** , and **reason** fields instead of the OpenAI shape. The gateway's check order — API key, model access, usage limits, input guardrails, routing/provider, output guardrails — determines which error a caller sees when multiple checks could fail, and 422 guardrail blocks on the way out still count tokens toward usage limits.

Attach a correlation ID to every API call so your operator can pinpoint the exact request behind a 422 guardrail block or 429 limit error in gateway logs.

```
$ curl https://your-gateway.example.com/v1/responses \
  -H "Authorization: Bearer sk-olga-..." \
  -H "X-Request-ID: checkout-7f3a2c" \
  -H "Content-Type: application/json" \
  -d '{"model": "gpt-4o-mini", "input": "Hello!"}'
```

– Full mechanism and exact codes given; header usage shown with runnable example.

product docs

02 CLI test deletion command NEW

70

Adds the `openlayer tests delete` CLI command, which archives tests by default; passing `--force` permanently deletes them.

Delete a specific test permanently from the CLI, bypassing the default archive behavior.

```
$ openlayer tests delete <test-id> --force
```

– Exact command and flag given with a runnable example.

Governance lifecycle stages, bias and f...

03 New Test Types NEW

70

Adds four new test types: Bias and fairness tests, with four test types checking protected groups and their intersections for outcome, score, performance, and metric disparities, plus an Analysis tab for charts, drill-down, and CSV export; Prompt injection tests, where you choose the detector and a sensitivity level (multilingual included), see why a detection fired, and pick the judge model; Freshness and Column freshness tests, which alert when a data source stops receiving data or a datetime column gets too old, with manual or automatic thresholds; and an `LLM-as-a-judge` test that scores every row in the window and explains its score on demand.

:

– Test types and UI location named; no config or API surface shown. Configurable options named but no API or config key shown. Threshold modes named; no exact config surface given. Thin description with no config or judge-selection detail.

Governance lifecycle stages, bias and f...

04 Per-step evaluation scope NEW

65

Adds the `evaluationScope: step` config field so tests can run on each individual trace step instead of only the final output.

Enable per-step evaluation so tests run against each trace step independently, not just the final output.

```
json
{
  "evaluationScope": "step"
}
```

– Named config field and file shown; no detail on which test types support it.

Governance lifecycle stages, bias and f...

05

Custom columns via JSONPath

NEW

60

Adds Custom columns — a named JSONPath into structured outputs, defined once per project and usable in tests, filters, and tables across every data source and development commit.

– Mechanism and scope named; no exact syntax example.

Governance lifecycle stages, bias and f...

THINNER COVERAGE BELOW

06

openlayer init CLI enhancements

IMPROVED

55

Extends `openlayer init` with support for Copilot CLI and Gemini CLI, custom CA bundles, a 'debug with an agent' option, and `.openlayerignore` for smaller uploads.

– Lists named options but gives no usage example.

Governance lifecycle stages, bias and f...

07

Governance registry and framework API endpoints

NEW

40

Adds governance registry and framework endpoints to the public API and SDKs, writable with a workspace API key.

– No endpoint paths or SDK methods named. Governance lifecycle stages, bias and f...

08

API key lifecycle management

NEW

40

Adds API key expiry dates, automatic rotation, and full API key management via the API.

– No endpoint names or parameters given for key management.

Governance lifecycle stages, bias and f...

↳ ALSO FROM THESE RELEASES

.

WAS THIS USEFUL?



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

GitHub's biggest shipment this window is a public-preview package firewall that intercepts npm downloads for CI/CD runners and coding agents, alongside a new `version_requirement` field in the Dependabot alerts API, one-click migration of branch protections to rulesets, several secret scanning updates, and a round of Copilot CLI additions covering MCP integration, tool controls, agent skills, and enterprise admin.

← WHAT SHIPPED • 8 FEATURES most completely described first ? what's the number?

01 Dependency version field in Dependabot alerts API NEW

95

Dependabot alert REST API responses add a `version_requirement` field to `dependency` objects, returning the resolved version (from lockfiles or the Dependency submission API) or declared constraint (for manifests without a lockfile) — eliminating the need to parse lockfiles or make a separate GraphQL `vulnerableRequirements` request. The field appears in individual alert responses, repository/organization/enterprise alert list responses, and `dependabot_alert` webhook payloads.

Pull a single Dependabot alert and inspect the new `version_requirement` field to see exactly which version of a vulnerable dependency is in your repo — no lockfile parsing needed.

```
$ curl -H 'Authorization: Bearer <token>' \
  'https://api.github.com/repos/<owner>/<repo>/dependabot/alerts/<alert_number>' \
  | jq '.dependency.version_requirement'
```

List all open Dependabot alerts for a repo and export each dependency name alongside its recorded version constraint for triage or SIEM ingestion.

```
$ curl -H 'Authorization: Bearer <token>' \
  'https://api.github.com/repos/<owner>/<repo>/dependabot/alerts?
state=open&per_page=100' \
  | jq '.[] | {package: .dependency.package.name,
version_requirement: .dependency.version_requirement, severity:
.security_advisory.severity}'
```

– Runnable API examples and exact field/endpoint scope given

Dependabot alerts REST API now includes...

02 Package firewall for npm in CI/CD and coding agents NEW

90

GitHub package firewall (public preview, GitHub Enterprise Cloud) sits between endpoints like Actions runners and coding agents and upstream package registries, inspecting metadata requests and package downloads. It checks packages against the GitHub Advisory Database for malware or vulnerability alerts, performs supply-chain checks including attestation downgrades, provides an exception request workflow for erroneously blocked packages consistent with other GitHub Advanced Security flows, and currently supports npm packages.

Route an Actions runner's npm traffic through the package firewall instead of pointing directly at the npm registry — blocks malicious packages at install time on CI/CD hosts.

📌 In your GitHub Enterprise Cloud admin console, go to Settings > Package firewall, configure the cooldown period and advisory/attestation policies, then follow [Protecting endpoints with package firewall](#) to point your runners at the firewall proxy instead of `registry.npmjs.org`.

– Admin console path and policy configuration steps given

GitHub package firewall in public preview...

03 Bulk migration of branch protections to rulesets NEW

80

Repositories can bulk-convert all branch protection rules to rulesets in a single action via Settings → Rulesets → [Import branch protection rules](#), with a preview of rulesets to be created, selection of enforcement status (Active, Evaluate, or Disabled), and an option to keep or delete the original branch protection rules after conversion.

Migrate all branch protection rules to rulesets at once, previewing what will be created before committing.

📌 1. Go to your repository's Settings → Rulesets. 2. Click [Import branch protection rules](#). 3. In the [Review import](#) dialog, preview rulesets to be created, check which

rules are already covered, and note any branch protection rules that produce multiple rulesets. 4. Choose an enforcement status: **Active**, **Evaluate**, or **Disabled**, then select which rulesets to create. 5. Click **Create rulesets**, then choose whether to keep or delete the original branch protection rules.

– Exact UI path and step-by-step workflow given [Bulk migrate branch protections to rule...](#)

04 Audit log business field switches to enterprise slug

BREAKING

80

Starting November 12, 2026, the **business** field in **secret_scanning_alert.*** and **secret_scanning_push_protection.bypass** audit log events changes from the enterprise display name to the enterprise slug; integrations filtering by display name must be updated to accept both formats or switch to **business_id** before rollout.

– Exact field names and migration deadline given, no code sample

[Updates to GitHub secret scanning \(2026...](#)

05 Exemption request number in secret scanning webhooks

NEW

65

Dismissal-request webhook payloads now include **exemption_request.number**, the repository-scoped request number also available via the REST API, enabling correlation of bypass requests across integrations.

– Named field but no endpoint or sample payload given

[Updates to GitHub secret scanning \(2026...](#)

06 Base64-encoded token revocation expands token types

IMPROVED

60

Base64-encoded token revocation (public preview), available in private and internal repositories on github.com, now supports encoded personal access tokens (classic), fine-grained personal access tokens, OAuth access tokens, and GitHub App user access tokens, with GitHub automatically decoding the value.

– Lists supported token types but no config or command to enable

[Updates to GitHub secret scanning \(2026...](#)

THINNER COVERAGE BELOW

07 Custom pattern API returns 422 for invalid requests

IMPROVED

50

The custom pattern API now returns **422** instead of **500** for requests rejected due to invalid arguments, letting integrations distinguish input errors from server failures.

– Named status codes but no endpoint path or example call

[Updates to GitHub secret scanning \(2026...](#)

Copilot CLI adds MCP (Model Context Protocol) server integration for connecting additional tools and context, allow/deny controls to restrict which tools Copilot CLI can invoke, agent skills for repeatable workflows with task-specific instructions, scripts, and resources, and enterprise-level administration controls.

– No commands, flags, or UI paths given for any addition [product docs](#)

NEW SUPPORTED PLATFORM [npm 12, pnpm 11, pnpm 12, and Bun 1.4 support in Dependabot](#)

↳ **BREAKING ON UPGRADE**

! On November 12, 2026, the `business` field in `secret_scanning_alert.*` and `secret_scanning_push_protection.bypass` audit log events will change from the enterprise display name to the enterprise slug. Integrations that filter by display name must be updated before rollout to accept both formats or switch to `business_id`.

Dependabot now supports npm 12, pnpm 11 and 12, and Bun 1.4 (2026-10-20)

[BLOG](#)

[STABLE](#)

WAS THIS USEFUL?



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

GitLab shipped an AI Governance Dashboard (beta) for auditing GitLab Duo and external coding agents, alongside an early-access Business Logic Security Analyzer that catches authorization and mass-assignment flaws traditional SAST scanners miss.

WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 AI Governance Dashboard for agent activity NEW 95

GitLab introduces the AI Governance Dashboard (beta), accessible at AI > Governance > Dashboard tab in a top-level group, giving a unified view of GitLab Duo Agent Platform agents and externally connected agents (Claude Code, Cursor, Codex). The `glab govern setup` subcommand (requires `glab` CLI 1.122.0+) installs Claude Code hooks to capture sessions, and `glab govern setup --agents codex,cursor` captures Cursor and Codex sessions via a background job that sends sessions every 30 minutes; `glab govern doctor` verifies that agent-capture setup is working correctly. A new `read_agent_artifacts` custom role ability grants dashboard access without requiring the Owner role on a top-level group, and a `Store AI audit events` toggle (in Settings > GitLab Duo > Change configuration > Data and privacy) enables persistent recording of every agent action as an audit event, which can be streamed to a SIEM system or downloaded as an immutable session copy for evidence.

.

Capture Claude Code, Cursor, and Codex sessions so all external agent activity shows up in the AI Governance Dashboard.

```
$ glab govern setup --agents codex,cursor
```

Confirm that agent-capture hooks are installed and sending sessions correctly before handing the dashboard to an auditor.

```
$ glab govern doctor
```

Enable AI audit event storage so agent actions start being recorded — without this, the dashboard shows no history.

📌 In your top-level group, go to Settings > GitLab Duo, select **Change configuration**, then enable **Store AI audit events** under Data and privacy. Requires Owner or Security Manager role.

– Names every CLI subcommand, custom role and config toggle

AI Governance: One view of agent activi...

02 Business Logic Security Analyzer in Duo Agent Platform NEW

88

GitLab Duo Agent Platform adds a Business Logic Security Analyzer (early access, GitLab Ultimate on [GitLab.com](https://gitlab.com)) that runs whole-repository SAST scans to catch broken object-level authorization (CWE-639), missing or improper authorization, and mass assignment — vulnerability classes pattern-based scanners cannot detect. Scans can be triggered via GitLab Duo Chat or the API, or run automatically on merge requests after pipelines pass by attaching a Business Logic scan profile to a project; findings report CWE ID, severity (e.g. Critical), and file location in the Vulnerability Report and the merge request security widget. The analyzer is backed by Anthropic's Claude Mythos 5.1 (a first for any GitLab capability) plus Claude Opus 5.5, Sonnet 5.5, Sonnet 5, and Sonnet 4.6 as GitLab-managed models through GitLab's AI gateway, with each scan consuming GitLab Credits scaled to the application's attack surface.

– Names CWE classes, models and credits but no exact command

Uncover the business logic vulnerabilit...

WAS THIS USEFUL?



RANK

Snyk AppRisk is a security tool that aggregates application and cloud data to prioritize and manage risks across an organization's software development lifecycle.

Snyk AppRisk adds configurable Secrets scan scheduling at the repository and bulk level, plus new Secrets coverage visibility in the Inventory All Assets view.

←→ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Secrets scan scheduling via Repo Monitor Configuration NEW

75

Adds **Repo Monitor Configuration** under **Inventory -> repository -> Settings** to set Secrets scan frequency to **Daily**, **Weekly**, or **Never** for a single repository, plus a bulk configuration option via **Inventory -> Configure Monitoring** to apply the same frequencies across many repositories at once. **Repo Monitor Configuration** is available as an Early Access feature enabled via Snyk Preview, and acts as a prerequisite for Secrets test configurations.

Reduce noise in low-risk repos by scheduling Secrets scans weekly rather than on every push.

📌 In the console, go to **Inventory -> repository -> Settings**, select a repository, open **Repo Monitor Configuration**, and set Secrets scan frequency to **Weekly**.

Disable Secrets scanning across all repositories in a bulk operation to save scan quota while a policy is being defined.

📌 In the console, go to **Inventory -> Configure Monitoring**, select all repositories, and set Secrets frequency to **Never**.

– Names exact UI paths and frequency options with usage steps.

2026-10-19-60816266

THINNER COVERAGE BELOW

02 Secrets coverage in Inventory All Assets view NEW

45

Adds Secrets coverage to the All Assets view in Inventory, showing each repository's Secrets coverage control, the associated Snyk Projects with Secrets results, and the Secrets issues found.

– Describes what the view shows but no navigation steps or example given.

2026-10-19-60816266

WAS THIS USEFUL?



DataDome

[Feedback](#)[Android SDK 3.1.0 \(2026-10-21\)](#)[BLOG](#)[RANK](#)

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

DataDome's Android SDK 3.1.0 adds native support for Android's predictive back gesture, alongside a minimum SDK/compileSdk bump required to use it.

↳ [WHAT SHIPPED](#) • [2 FEATURES](#) most completely described first [? what's the number?](#)

01 Predictive back gesture support NEW 86

The SDK now supports Android's predictive back gesture via a new `BackBehaviour` enum with values `BLOCKED`, `GO_BACKGROUND`, and `GO_BACK`, which is honoured on Android 13+ when the app sets `android:enableOnBackPressedCallback="true"` in its manifest or targets API 36+.

Opt in to Android predictive back so DataDome's BackBehaviour controls what happens when a user swipes back on a protected screen.

```
xml
<application
    android:enableOnBackPressedCallback="true"
    ... />
```

– Names enum values, manifest flag, and OS version gating [Android SDK 3.1.0 \(2026-10-21\)](#)

02 Raised minimum build requirements for predictive back 65

BREAKING

To use the new predictive back behaviour, apps must now depend on `androidx.activity` 1.6.1 and compile with `compileSdk` 33 or higher; projects on lower compile SDK versions will fail to build.

– Exact dependency and compileSdk versions given, no migration steps

[Android SDK 3.1.0 \(2026-10-21\)](#)

↳ **BREAKING ON UPGRADE**

! The SDK now requires `androidx.activity` 1.6.1 and apps must compile with `compileSdk` 33 or higher; projects on lower compile SDK versions will fail to build.

WAS THIS USEFUL?  

Fastly

Feedback



Fastly Terraform Provider (beta) 0.2.4 (2026-10-20)

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 Terraform Provider (beta) 0.2.4 extends its beta testing guide with a resource family workflow for discovering resources via Terraform's query capability.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Resource family discovery workflow in Terraform provider guide

IMPROVED

55

The beta testing guide now documents an explicit resource family workflow using `terraform query` combined with service-scoped discovery, helping users find resources tied to a given Fastly service.

- Names exact command but lacks mechanism detail Fastly Terraform Provider (beta) 0.2.4 ...

WAS THIS USEFUL?  

infisical

Feedback



v0.166.9

NOTES ↗

RANK

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

Infisical v0.166.9 adds AWS KMS-backed signing for internal CAs, S3-compatible storage for PAM session recordings, and broadens Agent Vault variable scoping, alongside UI revamps for findings and data stores, new audit-log-stream schema support, and name-suggestion flows across six creation surfaces.

↳ WHAT SHIPPED • 9 FEATURES most completely described first ? what's the number?

01 Context-based name suggestions across creation flows NEW

60

Infisical now suggests names from context during creation for secret syncs, app connections, honey tokens, certificate authorities, approval policies, and certificate policies.

– Names all six surfaces but no exact UI flow given v0.166.9

THINNER COVERAGE BELOW

02 AWS KMS signing for internal CAs NEW

55

The PKI module now supports AWS KMS-backed signing keys for internal certificate authorities, enabling hardware-protected certificate issuance.

– Names module and mechanism but no setup steps given v0.166.9

03 Audit log stream schema support NEW

55

Adds an audit_log_streams table to the sanitized database schema to support audit log streaming.

– Names exact table but lacks usage or endpoint detail v0.166.9

04 Automatic trial flow from pricing page IMPROVED

50

Selecting a plan on the pricing page now opens that plan's trial automatically after signup, and the pricing trial prompt is also shown to users who are already signed in.

– Describes the flow with a clear starting point v0.166.9

05 S3-backed PAM session recordings NEW

45

PAM session recordings can now be stored using an S3-compatible storage backend.

– Describes capability but no config detail provided v0.166.9

- 06

Offline license usage on billing page NEW

40
- The unified billing page now shows offline license usage.
- Names exact UI location, limited further detail v0.166.9
-
- 07

Findings view redesign IMPROVED

25
- The findings UI has been revamped with a redesigned findings view.
- Thin description, only notes a visual redesign v0.166.9
-
- 08

Data store and data source UI redesign IMPROVED

25
- The data store list has been revamped and the data source creation flow redesigned.
- Thin description, only notes a visual redesign v0.166.9
-
- 09

Breadcrumbs in top navigation NEW

25
- Adds page breadcrumbs to the top navigation on every route.
- Simple UI addition, no further mechanism described v0.166.9
-

WAS THIS USEFUL?



GitGuardian Self-Hosted

[Feedback](#)[2026.10 - Required \(2026-10-20\)](#)[BLOG »](#)

GitGuardian Self-Hosted is an on-premises version of GitGuardian's platform that scans code repositories and CI/CD pipelines to detect leaked secrets and credentials.

Adds Wiz sync, source-status filtering, auto-pausing/resuming unreachable integrations, a `beat.metricsPort` Helm value for scheduler metrics, new public API endpoints for presence checks and GitHub check-run management, and expanded audit log CSV exports.

Adds `beat.metricsPort` Helm value to configure the port on which the task scheduler exposes Prometheus metrics, with automatic `ServiceMonitor` creation when ServiceMonitors are enabled.

NEW INTEGRATION Adds Wiz integration that syncs active secret incidents from version control sources to Wiz daily as native secret findings, including detection context and a link back to GitGuardian but never the secret value — requires outbound HTTPS to your Wiz tenant.

Adds source-status visibility: every source now displays one of `Monitored`, `Unmonitored`, `Unreachable`, `Archived`, or `Deleted` states, filterable from the perimeter page and incident list, with an `is all of` operator, and saveable as a named view.

Adds public API endpoint to trigger a presence check on a secret.

Adds public API endpoints to find the GitHub check runs that reported an incident and to skip a blocking check run, enabling custom approval flows to unblock pull requests.

[+10 more](#)

NEW DETECTION [10 new detectors in Secrets Detection Engine v2.172](#)

UPDATED DETECTION [10 improved detectors in Secrets Detection Engine v2.172](#)

NEW DETECTION [2 new checkers with custom host support in v2.172](#)

NEW DETECTION [8 new analyzers in v2.172](#)

UPDATED DETECTION [1 analyzer update in v2.172](#)

NEW DETECTION [9 new detectors in Secrets Detection Engine v2.173](#)

UPDATED DETECTION [5 improved detectors in v2.173](#)

NEW DETECTION [1 new analyzer in v2.173](#)

UPDATED INTEGRATION [AI Hooks extended to Amazon Kiro and Junie CLI](#)

NEW SUPPORTED PLATFORM [Kubernetes 1.36 support](#)

↳ **BREAKING ON UPGRADE**

- ! Legacy analytics are removed on upgrade to 2026.10; if in-app analytics are disabled on your instance you must re-enable them before upgrading or your analytics dashboards will be lost.
- ! Embedded cluster preflight checks now require at least 12 CPU cores and 48 GiB of memory (up from 8 cores and 30 GiB); a smaller VM will fail the check.
- ! An upgrade to 2026.10 is refused with 'Not all required migrations have been applied' if background migrations from a prior release are still running; certain migrations must also complete on 2026.10 before upgrading to 2026.11.
- ! This is a required release and cannot be skipped.

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

YARA-X

[Feedback](#)

v1.22.0

[NOTES ↗](#)[RANK](#)

A rewrite of YARA in Rust.

YARA-X v1.22.0 adds 64-bit integer support to the rule language, a new PE offset-to-RVA helper, xor_key exposure in the C and Go APIs, and zip module visibility in yr dump, alongside Language Server and compiler diagnostic improvements.

↳ [WHAT SHIPPED](#) · [9 FEATURES](#) most completely described first [? what's the number?](#)

01 64-bit integer built-in functions NEW 65

Adds `int64`, `int64be`, `uint64`, and `uint64be` built-in functions, with support for unsigned 64-bit integers and integer literals greater than `i64::MAX`.

– Names exact functions but no usage example given. v1.22.0

02 pe.offset_to_rva function NEW 65

Adds `pe.offset_to_rva` to the `pe` module for converting file offsets to RVAs, useful for correlating raw file offsets with virtual memory layout.

Convert a PE file offset to its RVA inside a YARA-X rule condition, useful when correlating raw file offsets with virtual memory layout.

```
$ rule offset_example {
  condition:
    pe.offset_to_rva(0x400) > 0
}
```

– Named function with a runnable rule-condition example. v1.22.0

03 Language Server enhancements IMPROVED 65

Improves the Language Server with `Workspace Symbols` support, better `Code Completion` documentation, auto-importing modules on non-top-level symbol completion, and `--help` / `--version` CLI flags.

– Lists several concrete LSP additions including CLI flags. v1.22.0

THINNER COVERAGE BELOW

04 zip module support in yr dump NEW 50

Allows the `yr dump` command to show information from the `zip` module.

– Names command and module; no example output shown. v1.22.0

05 **xor_key exposed in C and Go APIs** NEW 45

Exposes `xor_key` for matching patterns in the C and Go APIs.

– Names the field and APIs but no usage detail. v1.22.0

06 **Compiler diagnostic warnings and documentation links** IMPROVED 35

Raises warnings on implicit sign conversions that change constant values and on unary operations on unsigned integers, and links compiler errors and warnings to their documentation pages.

– Describes new diagnostics but no example or mechanism detail. v1.22.0

07 **Constant folding for division and modulus** IMPROVED 25

Implements constant folding for integer division and modulus operations.

– Bare statement of a compiler optimization, no detail. v1.22.0

08 **Protobuf integer field width and signedness preserved** IMPROVED 25

Keeps the width and signedness of protobuf integer fields.

– Thin description with no mechanism or scope. v1.22.0

09 **Performance and memory improvements** IMPROVED 5

Delivers performance and memory consumption improvements.

– No numbers, scope, or mechanism given. v1.22.0

WAS THIS USEFUL?  

CLOUD & CONTAINER SECURITY

Chainguard

Feedback



changelog · Oct 21, 2026

CHANGELOG ↗

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

Chainguard ships six new products in beta/EA: Custom Assembly Overlays, Harness (AI runtime vuln finder), Checks (hardened CI microVMs), Workspaces (isolated agent microVMs), and Libraries for .NET and Go.

Adds the `image_overlay` Terraform resource to manage Custom Assembly Overlays and their bindings, enabling per-tag package and configuration customization within a multi-version repository.

Adds Custom Assembly Overlays (Early Access) to apply packages, environment variables, annotations, certificates, and user accounts to exact tags, `-dev` variants, or entire organizations without breaking unrelated versions.

Adds version templates (e.g. `py{{major}}.{{minor}}-cryptography`) so each tag in an overlay binding resolves its own concrete package version automatically.

NEW INTEGRATION Adds Chainguard Harness (Beta), an AI-powered runtime vulnerability finder that surfaces only exploitable findings reproducible in a running application and returns engineer-reviewed fixes.

Adds Chainguard Checks (Beta), replacing standard CI runners with fast, hardened microVMs that enforce egress policies by default and allow engineers or agents to run tests before pull request submission.

+4 more

NEW SUPPORTED PLATFORM

New Chainguard container images — week of 2026-10-19

NEW SUPPORTED PLATFORM

New Chainguard container images — week of 2026-10-12

⚠️ BREAKING ON UPGRADE

- ! After December 15, 2026, Chainguard will sign new artifacts with Sigstore v3 only; any tooling still running Cosign v2 or v3.1.0 and earlier will be unable to verify those artifacts. Upgrade to `cosign v3.1.1` or later before that date.
- ! Custom verification tooling built against the Cosign v2 libraries must be reviewed and updated against the v3 library API, which changed in Cosign v3.
- ! Registries used to mirror Chainguard artifacts must support the OCI 1.1 Referrers API; registries that do not support it may fail to copy v3 signatures alongside mirrored artifacts.

WAS THIS USEFUL?



kata-containers

[Feedback](#)

4.3.0

[NOTES ↗](#)[RANK](#)

Kata Containers is an open source project and community working to build a standard implementation of lightweight Virtual Machines (VMs) that feel and perform like containers, but provide the workload isolation and security advantages of VMs. <https://katacontainers.io/>

Kata Containers 4.3.0 adds rootless SEV-SNP and TDX confidential VMs, flexible NVIDIA driver sourcing for GPU workloads, and expands guest-extensions beyond QEMU, alongside a breaking removal of the `image_guest_pull` driver option.

←→ WHAT SHIPPED • 8 FEATURES most completely described first ? what's the number?

01 Custom NVIDIA driver sources via `driver.url` in `nvgpu` NEW

80

Adds support for a dedicated driver userspace repository for NVIDIA GPUs via `driver.url` in `nvgpu`, and accepts a complete driver source tarball as `driver.url` for NVIDIA GPU builds.

Point `nvgpu` at a dedicated driver userspace repository or a full driver source tarball so every GPU runtime class component is built and shipped.

```
yaml

# In the nvgpu driver configuration, set driver.url to a
# dedicated repository or a complete tarball:
driver:
  url: "https://internal-mirror.example.com/nvidia-driver-
  570.tar.gz"
```

– Exact config key with runnable example configuration 4.3.0

02 Removal of `image\_guest\_pull` driver option metadata BREAKING

70

The `image_guest_pull` driver option metadata is removed from `genpolicy`, `agent`, and `runtime(-rs)` – configurations or policies that set this field will break.

– Clear migration note naming affected components 4.3.0

03 kata-deploy device provisioning and configuration controls

IMPROVED

65

Adds kata-deploy support to provision host devices for rootless VMMs and to detect external Node Feature Discovery (NFD) by label. Also restricts containerd

`pod_annotations` to the per-installation configuration and honors `imagePullSecrets` on `kubecttl-image` hook Jobs.

– Names `pod_annotations` and `imagePullSecrets`, a clear starting point 4.3.0

04 **Guest-extensions support expanded beyond QEMU** IMPROVED

60

Expands guest-extensions support to VMMs beyond QEMU, updates Cloud Hypervisor to v53.0, and bumps the guest kernel to 6.18.52.

– Specific version numbers but no usage steps 4.3.0

05 **New agent capabilities: secret unsealing, cgroup throttling, fsmount fallback**

NEW

60

Adds agent capabilities to unseal sealed secret environment variables during exec requests, report cgroup v2 `CPU` throttling time, and fall back to `fsmount` when overlay options exceed 4K.

– Names `fsmount` and 4K threshold but no direct invocation shown 4.3.0

THINNER COVERAGE BELOW

06 **genpolicy image policy verification improvements** IMPROVED

55

Verifies guest-pulled image references in `genpolicy`, adds a default `PATH` when the container image environment lacks one, and names the missing Secret or ConfigMap when an env source is absent.

– Describes behaviour changes but no command or flag 4.3.0

07 **Rootless SEV-SNP and TDX confidential VMs** NEW

50

Enables rootless SEV-SNP and TDX in `runtime-rs` /QEMU, allowing confidential VMs to run without root privileges.

– Names SEV-SNP, TDX, `runtime-rs` but no concrete setup steps 4.3.0

08 **Minimum memory limit enforcement in kata-webhook** NEW

25

Adds ability to enforce a minimum memory limit in the kata-webhook.

– Bare description, no mechanism or config key given 4.3.0

↳ **BREAKING ON UPGRADE**

! The `image_guest_pull` driver option metadata is removed from `genpolicy`, `agent`, and `runtime(-rs)` – configurations or policies that set this field will break.

WAS THIS USEFUL?



Google distroless

Feedback



changes since 2026-10-17

CODE ↗

Language focused docker images, minus the operating system.

Distroless blocks compromised CA roots and switches Node.js images to system CA trust store

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-17
https://github.com/GoogleContainerTools/distroless.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-17
```

Blocks two compromised CA roots ([AAA Certificate Services](#) and 'emSign ECC Root CA - G3') from the `ca-certificates.crt` bundle, matched by SHA-256 SPKI hash to catch renames and re-issues.

Adds a `TestBlockedRoots` test in `common/cacerts_test.go` that validates every image archive against a `blockedSPKIs` denylist, covering both the PEM bundle at `etc/ssl/certs/ca-certificates.crt` and the Java trust store at `etc/ssl/certs/java/cacerts`.

Allows skipping the mandatory 1-day soak time for image updates, enabling faster emergency pushes.

WAS THIS USEFUL?



Sysdig Secure

[Feedback](#)[changelog · Oct 21, 2026](#)[CHANGELOG ↗](#)

RANK

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

Sysdig Secure reworks SIEM integration filtering, replacing the old blanket data-forwarding setting with granular, rule-based filters.

↳ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

01 SIEM forwarding rules replace Data to Send field

BREAKING

81

Adds SIEM forwarding rules to SIEM & Data Platforms integrations, replacing the **Data to Send** field: you can attach one or more rules to each integration specifying a data type (e.g., Runtime Policy Events, Activity Audit) plus optional filters such as severity or policy name for Runtime Policy Events, or type for Activity Audit, so only matching data is forwarded. A new **Rules** column on the integrations list shows the number of forwarding rules configured per integration. Existing integrations that relied on the old **Data to Send** field must be reconfigured using the new rules model.

⋮

Reduce SIEM noise by forwarding only high-severity Runtime Policy Events to your SIEM integration, rather than all data.

📌 In Sysdig Secure, go to Integrations > SIEM & Data Platforms, select your integration, then under Forwarding Rules add a rule: set data type to **Runtime Policy Events** and filter by severity **High** or **Critical**. Save the integration.

– Names exact field, column, and filter types; includes migration note.

[changelog-20261021-d6245d2b](#)

↳ BREAKING ON UPGRADE

! The **Data to Send** field in SIEM & Data Platforms integrations is replaced by forwarding rules; existing integrations relying on that field will need to be reconfigured with the new rules model.

WAS THIS USEFUL?



Plexicus

[Feedback](#)

Plexicus v1.2.55: C# AI-SAST, PlexiFlow Editor Rebuild & Dashboard KPI Accuracy

NOTES >

Plexicus validates security findings with evidence and helps your team fix what matters. AI Swarm Pentest, Deep Code Analysis, Remediation and XBOM.

PlexiFlow automation editor rebuilt with undo, visual conditions, searchable pickers, and a schedule builder; Dashboard KPIs verified accurate against live data.

Adds undo functionality to the PlexiFlow automation editor.

Adds visual conditions to the PlexiFlow automation editor for building conditional logic without hand-editing.

Adds searchable pickers to the PlexiFlow automation editor for faster node selection.

Adds a schedule builder to the PlexiFlow automation editor for time-based automation triggers.

Verifies and corrects Dashboard KPI accuracy against live data.

NEW SUPPORTED PLATFORM

C# and .NET AI SAST coverage

WAS THIS USEFUL?



Rapid7 InsightCloudSec

[Feedback](#)[changelog · Oct 21, 2026](#)[DOCS](#)[BLOG](#)[RANK](#)

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

InsightCloudSec's October release adds new Insights and a Query Filter for unprotected database security groups, unencrypted DynamoDB tables, and exposed or unencrypted cache clusters, alongside an expanded RDS harvester and refined compliance mappings for two existing checks.

← WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Cache cluster security Insights NEW 70

Adds [Cache Instance Security Group Allows Ingress From the Public](#) Insight to identify cache instances whose security groups permit public ingress, mapped to CIS Controls v8.1.2 safeguard 12.2, CMMC Level 1 and 2, NIST SP 800-53 Rev. 5, NIST SP 800-171, and NIST CSF v2.0. Also adds [Cache Database Cluster Transit Encryption Disabled](#) Insight to flag Cache Database Clusters without Transit Encryption, mapped to CIS AWS Database Services Benchmark control 6.2 and CIS Controls v8.1.2 safeguard 3.10.

– Two named Insights with full compliance mapping detail [2026-10-21-f5946d32](#)

02 Org-shared AMI exclusion in image ownership filter IMPROVED 65

Updates the [Instance With Image Owned Outside Of Organization](#) Query Filter so instances running AMIs shared from accounts within your organization are no longer flagged, reducing false positives in custom Insights using this filter.

– Clear before/after behavior change on a named filter [2026-10-21-f5946d32](#)

03 DynamoDB default encryption key Insight NEW 60

Adds the [Distributed Table Using Provider Default Encryption Key \(DynamoDB\)](#) Insight to detect DynamoDB tables using the AWS-owned default key instead of an AWS-managed or customer-managed KMS key.

– Clear mechanism and named Insight, but no config/API surface [2026-10-21-f5946d32](#)

04 Expanded RDS harvester security group evaluation IMPROVED 60

Expands the Amazon RDS harvester to evaluate VPC Security Group configurations across RDS DB Instances and Clusters, including Aurora, using existing harvested data.

– Describes scope and mechanism but no direct user action 2026-10-21-f5946d32

05 **RDS encryption Insight compliance mapping update** IMPROVED

 60

Updates the [Database Instance not Encrypted \(RDS\)](#) Insight to support CIS Amazon Web Services Database Services Benchmark v2.0.0 Control 2.2 and adds compliance mappings for NIST SP 800-53 Rev. 5, NIST SP 800-171 Rev. 2, and CMMC Level 2.

– Names exact standards added but no new mechanism 2026-10-21-f5946d32

THINNER COVERAGE BELOW

06 **Database instance without security group detection** NEW

 53

Adds the [Database Instance Without Security Group](#) Query Filter to identify database instances and clusters lacking an attached security group, plus a matching [Database Instance Without Security Group](#) Insight that flags database instances and clusters with no attached security group configured.

– Names exact filter/Insight but no deeper mechanism given 2026-10-21-f5946d32

WAS THIS USEFUL?



Plerion

Feedback



Workload scanning covers Lambda web functions:

NOTES >

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

Lambda workload scanning now covers web function revisions with no new toggle required; unscanned functions surface the S3 bucket name for remediation.

Surfaces the S3 bucket name on any web function that cannot yet be scanned, guiding operators to share that bucket with Plerion the same way as for an AgentCore runtime deployed from a code zip.

UPDATED VULNERABILITY CHECK

Lambda workload scanning extended to Lambda web functions

WAS THIS USEFUL?



DETECTION & RESPONSE

Sophos Core Agent

Feedback



Version 2026.3.0.2090Windows 10 64-bit and later

NOTES ↗

RANK

Sophos Core Agent is the endpoint software component of Sophos Central that collects telemetry and enforces security policies such as malware protection and threat detection on managed devices.

Sophos Core Agent 2026.3.0.2090 consolidates its local user-facing tools into a single unified component.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Unified Sophos Endpoint UI component

IMPROVED

40

Sophos Endpoint UI and Sophos Self Help Tool are merged into a single unified component called Sophos Endpoint UI.

– Names the merged components but no behavioural detail

Version 2026.3.0.2090Windows 10 64-bit ...

WAS THIS USEFUL?



Sophos Central Server Core Agent

Feedback



Version 2026.3.0.2104Windows Server 2016 and later

NOTES

RANK

Sophos Central Server Core Agent is an endpoint security agent installed on servers that enables management, threat detection, and protection through the Sophos Central cloud console.

Sophos Central Server Core Agent merged its endpoint UI and self-help tool into a single unified component in version 2026.3.0.2104.

WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Unified Sophos Endpoint UI IMPROVED

30

The Sophos Endpoint UI and Sophos Self Help Tool have been merged into a single unified component called Sophos Endpoint UI, for Windows Server 2016 and later.

– Names the merged component but no mechanism or UI path detail.

Version 2026.3.0.2104Windows Server 201...

WAS THIS USEFUL?



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

Huntress shipped an on-demand Identity Posture Assessment for Microsoft 365 with API access and shareable PDF reports, and introduced Managed ITDR Lookback, which automatically scans 180 days of M365 login history at tenant onboarding to surface pre-existing compromises.

WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 Managed ITDR Lookback for M365 onboarding NEW 81

Lookback automatically runs Huntress Unwanted Access detections across up to 180 days of Microsoft 365 login history when an eligible tenant onboards, surfacing pre-existing compromises from day one. Findings appear as a dedicated section of the Identity Security Assessment, delivered by email and visible in the portal within 48 hours of onboarding; active compromises trigger SOC investigation, a critical Incident Report, and active session revocation, while previously remediated activity surfaces only as historical findings. It runs once, automatically, for eligible Microsoft 365 tenants (including trial tenants) onboarded on or after October 19, 2026, and is not available for GCC, GCC High, or CPV tenants.

When onboarding a new client to Managed ITDR, review the Lookback section of the Identity Security Assessment in the Huntress portal to see whether any pre-existing compromises were detected across the past 180 days of M365 login history.

In the Huntress portal, go to the newly onboarded tenant > Identity Security Assessment > Lookback section (available within 48 hours of onboarding). Review flagged historical findings; active compromises will already have a critical Incident Report opened by the SOC.

- Detailed mechanism, timing, exclusions, and portal path given

Find Out Who Was in Your Microsoft 365 ...

02 Identity Posture Assessment in ISPM NEW 70

Huntress introduces the Identity Posture Assessment, an on-demand report that evaluates a Microsoft 365 tenant against the Huntress Identity Security Framework, scoring only licensed controls and ranking gaps by severity. A new Huntress API endpoint lets partners retrieve the report programmatically for integration into existing workflows, and a co-branded, shareable PDF can be generated from the ISPM dashboard

via the Quick Assessment card to support client onboarding, QBRs, and prospect conversations.

Generate and download a co-branded PDF posture assessment for a client directly from the Huntress portal ISPM dashboard.

1. In the Huntress portal, navigate to the ISPM dashboard.
2. In the Quick Assessment card, select the target organization from the dropdown.
3. Click Generate to download the PDF report.

– Names API endpoint and dashboard card but not exact path

Now Live: Managed ISPM Identity Posture...

WAS THIS USEFUL?



Trend Vision One

Feedback



TrendAI Vision One™ IAM for CEGP (2026-10-20)

[BLOG](#)

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

Cloud Email Gateway Protection gains TrendAI Vision One™ IAM directory sync with a migration path from the legacy Directory Synchronization Tool via the IAM Connection tab.

NEW INTEGRATION

Adds TrendAI Vision One™ IAM as a directory synchronization backend for Cloud Email Gateway Protection, supporting groups, valid recipients, and email alias sync.

Migration eligibility and report available in the **IAM Connection** tab; switching triggers automatic cleanup of legacy Directory Synchronization Tool data and configuration.

WAS THIS USEFUL?



Tenzir

Feedback



Tenzir Test v2.3.1: Automatic package loading in Python tests (2026-10-21)

BLOG ↗

RANK

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

Tenzir Test v2.3.1 brings Python tests in line with TQL tests by automatically loading enclosing and configured packages.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Automatic package loading in Python tests IMPROVED

50

Python tests now automatically load their enclosing package and configured package directories, matching the behavior already available in TQL tests. Python executors now capture package settings at construction time, so explicit environment mappings reliably select packages.

– Explains mechanism but no concrete config keys or commands given

Tenzir Test v2.3.1: Automatic package l...

WAS THIS USEFUL?



Cortex

[Feedback](#)

4.2.0

[NOTES](#)[RANK](#)

Cortex: a Powerful Observable Analysis and Active Response Engine

Cortex 4.2.0 strengthens OAuth2 security with PKCE and adds two Kubernetes job configuration options for private registry images and job retry behavior.

←→ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 Kubernetes job configuration additions

NEW

76

Adds a `job.kubernetes.imagePullSecrets` configuration option so analyzer and responder job images can be pulled from a private registry, and sets `backoffLimit` to `0` for analyzer and responder jobs in Kubernetes.

Pull analyzer/responder job images from a private registry in Kubernetes by referencing an image pull secret.

```
job {
  kubernetes {
    imagePullSecrets = ["my-registry-secret"]
  }
}
```

– Names exact config key and example usage, but no deeper mechanism explained 4.2.0

THINNER COVERAGE BELOW

02 PKCE support for OAuth2

NEW

25

Cortex 4.2.0 adds PKCE support in OAuth2 authentication.

– Bare mention with no configuration detail or usage shown 4.2.0

WAS THIS USEFUL?



Shuffle

Feedback



v2.3.1-rc1

NOTES ↗

RANK

Shuffle: A general purpose security automation platform. Our focus is on collaboration and resource sharing.

Shuffle v2.3.1-rc1 improves reliability of AI Agent-driven workflow execution in hybrid deployments.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 AI Agent execution improvements in hybrid workers IMPROVED

30

Adds graceful worker stop for AI Agent actions in hybrid deployments, and improves app name resolution and workflow execution validation when the executor is an AI Agent.

– Describes changes in prose without concrete config or API detail. v2.3.1-rc1

WAS THIS USEFUL?



◆ Threat Intelligence

Cloudmersive

Feedback



Cloudmersive Private Cloud Fraud AI API v2-2-0 Released (2026-10-20)

BLOG ↗

Cloudmersive is a developer API platform offering cloud-based services for file scanning and malware detection, document conversion, image processing, and data validation.

NEW DETECTION

Additional document classes for fraud detection

WAS THIS USEFUL?



Databricks

[Feedback](#)

[Asana connector is available in Beta \(2026-10-20\)](#)
[BLOG](#)
[RANK](#)

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

Databricks this window added a new Lakeflow Connect source (Oracle Eloqua, Beta), per-user row filtering via OAuth identity claims, a restructured Lakebase Insights experience with on-demand runs, access control for Genie budgets, and GA file upload in Genie conversations.

←→ **WHAT SHIPPED** • **5 FEATURES** most completely described first [? what's the number?](#)

01 Lakebase Insights moves to per-project Settings with on-demand runs

94

BREAKING

Lakebase Insights configuration now lives in each project's **Settings** (for users with **CAN MANAGE**), removing the need for a separate observability configuration, and supports on-demand Insights runs at any time with or without a scheduled analysis frequency set in **Settings**, reading telemetry from **system.lakebase** system tables to diagnose issues and recommend next steps. Enabling the Insights preview also auto-enables **Lakebase Observability in System Tables** if not already active, and admins can enable that observability preview independently to capture telemetry without Insights. Breaking: jobs previously scheduled through an observability configuration have stopped and their schedules do not carry over — those configurations no longer write telemetry to their Unity Catalog destinations, so dashboards retain historical data but receive no new data; to resume scheduled insights, enable Insights and set an analysis frequency in each project's **Settings**.

Trigger an immediate, unscheduled Insights analysis on a specific project without waiting for a scheduled run.

📌 In the Databricks UI, open the project, go to Settings > Insights, and click **Run now**.

Enable Insights for a workspace so telemetry begins flowing into system.lakebase for all subsequent project analyses.

📌 In the Databricks UI, go to the workspace admin panel, open Previews, and enable **Lakebase Insights**.

– Names config location, system table, and exact migration steps.

Lakebase Insights moves to per-project ...

02 Oracle Eloqua connector in Lakeflow Connect (Beta)

NEW INTEGRATION

NEW

75

Adds a managed Oracle Eloqua connector to Lakeflow Connect in Beta for ingesting Eloqua data into Databricks. It supports OAuth user-to-machine (U2M) authentication, incremental ingestion, UI-based pipeline authoring, Declarative Automation Bundles, and API-based pipeline authoring; workspace admins can enable or disable the connector from the Previews page.

Enable the Oracle Eloqua connector Beta for your workspace before building an ingestion pipeline.

In the Databricks console, go to Settings > Previews, find **Oracle Eloqua connector**, and toggle it on.

– Names auth mode and authoring methods; enable path given but no code sample.

Oracle Eloqua connector is available in...

03 Access control lists for Genie budgets

NEW

75

Adds ACLs to Genie budgets, letting account and workspace admins grant **Budget Manager** or **Budget Viewer** roles to specific users, groups, or Databricks service principals — for example assigning **Budget Viewer** to a service principal via Account > Budgets > Permissions for read-only automated cost reporting.

Delegate read-only budget visibility to a service principal so automated reporting jobs can access cost data without admin rights.

In the Databricks console, go to your Account > Budgets, select a budget, open Permissions, and assign the **Budget Viewer** role to the desired user, group, or service principal.

– Names both roles and exact UI navigation to assign them.

Delegate budget access with access cont...

Delegate access to Genie budgets (2026-...

THINNER COVERAGE BELOW

04 Row filtering per end user via OAuth custom identity claims

NEW

55

Embeds custom identity claims in a service principal's OAuth token, usable in dynamic views or ABAC policies to filter rows per end user — enabling one service principal to serve many users while each sees only their own data.

– Mechanism described but no setup steps or policy syntax.

05 File upload in Genie One and Genie Agents (GA)

UPDATED COMPLIANCE FRAMEWORK

IMPROVED

45

File upload in Genie One and Genie Agents is now generally available, letting users analyze uploaded files alongside existing workspace data through natural-language conversations; it no longer requires a preview toggle from workspace admins.

– States GA status and scope but no usage steps or limits.

File upload in Genie One and Genie Agen...

NEW INTEGRATION

DeepSeek V4.1 Flash (1017) on Unity Gateway

↳ ALSO FROM THESE RELEASES

:

↳ BREAKING ON UPGRADE

! Jobs scheduled through an observability configuration have stopped and their schedules do not carry over; those configurations no longer write telemetry to their Unity Catalog destinations, so dashboards retain historical data but receive no new data. To resume scheduled insights, enable Insights and set an analysis frequency in each project's

[Settings](#).

Asana connector is available in Beta (2026-10-20)

BLOG ↗

STABLE

DeepSeek V4.1 Flash (1017) is now available on Unity Gateway (2026-10-19)

BLOG ↗

STABLE

File upload in Genie One and Genie Agents will soon be available by default for compli

BLOG ↗

STABLE

File upload in Genie One and Genie Agents — upload files to a conversation and analyze them in natural language alongside your other data.

Enables file upload in Genie One and Genie Agents by default for workspaces with the compliance security profile enabled and a supporting compliance standard, allowing natural-language analysis of uploaded files alongside other data.

WAS THIS USEFUL?



defectdojo

Feedback



3.4.201

NOTES ↗

RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

DefectDojo v3.4.201 ships Keyboard Triage for rapid finding status changes, CRA Article 14 breach-deadline reporting, and self-managed internal CA/certificate rotation via `certs rotate-internal`, alongside GA of Change Plans and the Needs Attention inbox, new import/asset-identity resolution flags, and connector health visibility.

↳ WHAT SHIPPED • 10 FEATURES most completely described first ? what's the number?

01 Internal CA and service certificate management

BREAKING

95

Adds a `certs rotate-internal` doctor command to replace shipped internal CA and service certificates, backing up old files first, plus a `certs status` doctor command and a start preflight warning for existing installs still using shipped copies. New installs now generate their own internal CA and service certificates instead of trusting the bundled copies, and container images no longer trust the emSign root CAs.

Rotate internal CA and service certificates on an existing self-hosted install that is still using the shipped bundle copies.

```
$ certs rotate-internal
```

– names exact commands, runnable example, and breaking trust-store change [product docs](#)

02 Configurable uwsgi startup probe

NEW

90

Adds `django.uwsgi.startupProbe.enabled` (on by default) as a TCP check on the uwsgi port allowing 5 minutes to start, with `django.uwsgi.startupProbe.path` available to revert to the earlier HTTP check.

Override the uwsgi startup probe from the default TCP check back to an HTTP check in a Helm-based deployment.

```
yaml

django:
  uwsgi:
    startupProbe:
      enabled: true
      path: /healthz
```

– exact config keys, defaults, and runnable override example [product docs](#)

03 Imports follow renamed or moved Assets NEW 85

Adds **Imports Follow Renames and Moves** feature flag (on by default) so imports that name an Asset or Organization by a former or moved-out-of name resolve to the current Asset instead of failing or creating a duplicate; disable the flag in Settings > Feature Flags to revert to the old behavior.

Disable import rename/move resolution if you want imports that reference old Asset names to fail rather than silently land in the renamed Asset.

 In the DefectDojo UI, go to Settings > Feature Flags and turn off **Imports Follow Renames and Moves**.

– names flag, default state, and exact toggle location [product docs](#)

04 Keyboard Triage mode for Findings NEW 70

Adds a Keyboard Triage mode — opened with the **Triage** button or the **t** key on any Findings list — to walk findings one at a time and set status from the keyboard.

– names trigger controls but no further scope detail [product docs](#)

05 CRA Article 14 reporting NEW 70

Adds CRA Article 14 reporting without requiring PSIRT: statutory deadlines warn and breach, deadline notifications reach users who can view the Asset, and reports can be started from the UI; the **New Severe Incident Report** action is renamed to **New CRA Report**.

– describes notification flow and renamed action but no exact path [product docs](#)

06 Change Plans and Needs Attention inbox reach GA IMPROVED

 70

Change Plans and the Needs Attention inbox are now out of beta and enabled on every instance, with their feature flags removed from the Feature Flags page; Change Plans gain warnings alongside errors, show the SLA configuration a criticality change maps to, and support discard or rebuild without operations on Assets still recalculating.

– full behavior change described but no exact UI path [product docs](#)

THINNER COVERAGE BELOW

07 Asset Identity controlled by feature flag BREAKING

 55

Adds **Asset Identity** feature flag (off by default) to control Asset Identity in place of the former deployment setting; instances that previously had the deployment setting set to **true** keep it enabled after upgrade.

– names flag and migration behavior but no toggle path given [product docs](#)

08 Connector health visibility NEW

 50

Adds a **Zero Results** chip on connector tiles and a Connector Health Warning when a connector's syncs keep importing nothing.

– names UI indicators but no threshold or path detail [product docs](#)

09 What's New sidebar button NEW

 45

Adds a **What's New** button at the bottom of the sidebar showing release highlights for the currently running version.

– thin UI description with no further mechanism [product docs](#)

10 Qualys and Coverity connector improvements UPDATED INTEGRATION

IMPROVED

 45

Qualys VMDR imports and the Qualys connector now record each host's name as a host Location, and the Coverity connector now accepts **http** Coverity Connect URLs.

– names two connectors and their changes but no deeper mechanism [product docs](#)

↳ ALSO FROM THESE RELEASES

.
. .
.

↳ BREAKING ON UPGRADE

! The **New Severe Incident Report** action is renamed to **New CRA Report**.

- ! The **Asset Identity** feature is now controlled by the **Asset Identity** feature flag instead of the former deployment setting; instances that set the deployment setting to **true** retain it, but the old setting name is no longer the control surface.
- ! Container images no longer trust the emSign root CAs (CA trust stores refreshed).
- ! New installs generate their own internal CA and service certificates instead of using the copies shipped in the bundle; existing installs still on the shipped copies receive a warning and must run **certs rotate-internal** to replace them.

[3.4.201](#)[NOTES ↗](#)[LATEST](#)[STABLE](#)

DefectDojo 3.4.201 reports the Qualys VMDR asset name as a host location for imported findings.

↳ **GET THIS VERSION**

```
$ git clone --branch 3.4.201 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.4.201
```

Reports the Qualys VMDR asset name as a host location for findings imported from Qualys VMDR.

WAS THIS USEFUL?



syzkaller

Feedback



changes since 2026-10-19

CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller gains Linux device-mapper and MD RAID syscall descriptions, plus bluetooth and USB fuzzing focus in syz-cluster

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-19
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-19
```

Adds `exclude_path_regexps` support in `syz-cluster` fuzz targets to exclude specific paths from fuzzing scope

Adds syscall descriptions for the Linux device mapper control interface (`sys/linux`)

Adds syscall descriptions for Linux MD software RAID interfaces (`sys/linux`)

Adds bluetooth fuzzing focus in `syz-cluster`

Expands enabled syscalls for USB focus in `syz-cluster`

+2 more

WAS THIS USEFUL?



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

AI tool-call limits are now configurable via env vars and API; LonTalk protocol added for asset discovery.

Adds `RUNZERO_LLM_MAX_TOOL_CALLS_PER_TURN` and `RUNZERO_LLM_MAX_TOOL_CALLS_PER_DEEP_TURN` environment variables for self-hosted operators to set platform-wide AI tool-call limits (defaults: 8 and 20).

Exposes AI tool-call limits for accounts and organizations using their own AI provider via the AI configuration pages and through the API.

NEW SUPPORTED PLATFORM Adds LonTalk protocol support for asset discovery.

Expands self-hosted console environment variable documentation to cover AI provider, MCP session, access control, screenshot, and TLS trust settings.

UPDATED DETECTION Fingerprint improvements for BrightSign, Mitel, Emerson DeltaV, OPC UA, Beckhoff Automation, Cisco, and Loytec

UPDATED INTEGRATION Vulnerability validation for Tenable products improved

WAS THIS USEFUL?



◆ Endpoint Management (UEM/RMM)

Iru

Feedback



Mac Custom Apps: Now with support for 30 GB installers

NOTES

RANK

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

Iru raised the installer size limit for Mac Custom App Library Items to 30 GB and made large downloads more resilient with background downloading and auto-resume.

←▶ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 30 GB installer support for Mac Custom Apps IMPROVED

55

Mac Custom App Library Items now support installers up to 30 GB. The Iru Agent downloads these in the background so other Library Items are not blocked, and interrupted downloads caused by network hiccups or devices sleeping now auto-resume.

– Names limit and behaviors but no UI steps or config

Mac Custom Apps: Now with support for 3...

WAS THIS USEFUL?



IDENTITY, ACCESS & DATA PROTECTION

Passbolt

Feedback



snapshot-20261021

NOTES ↗

RANK

Manage and share passwords securely with Passbolt. Open source, audited and built for teams that need collaboration, compliance and control.

Passbolt 5.17.0 introduces phishing-resistant passkey MFA and adds PIN code support to QuickAccess and Offline Mode, alongside smaller improvements to permission review and autofill detection.

↳ WHAT SHIPPED • 4 FEATURES most completely described first ? what's the number?

01 Passkey-based MFA option NEW 35

Passbolt adds phishing-resistant passkeys as a new multi-factor authentication option, giving users an alternative to existing MFA methods.

– Names the feature but no setup steps or mechanism snapshot-20261021

02 Confirmation step on resource permission moves IMPROVED 28

When resources are moved, Passbolt now confirms permission changes, surfacing a review step before the move completes.

– Describes the change but no UI path or config snapshot-20261021

03 PIN codes for QuickAccess and Offline Mode NEW 25

Extends both QuickAccess and Offline Mode to support PIN codes, per the 5.17.0 release summary.

– Mentioned only in summary, no detail on behavior snapshot-20261021

04 Improved autofill detection on modern websites IMPROVED 15

Improves autofill detection across modern websites, per the release notes.

– Bare statement with no mechanism or scope snapshot-20261021

↳ ALSO FROM THESE RELEASES

•
•
•





WAS THIS USEFUL?



Didit

[Feedback](#)[snapshot-20261021](#)[NOTES ↗](#)

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

Didit adds KYB data tiers, AML/monitoring APIs, digital-ID wallet production launch, Face Search toggle, iGaming TM presets, and Copilot undo/slash skills.

UPDATED INTEGRATION Adds `POST /v3/users/aml-monitoring/` and `POST /v3/businesses/aml-monitoring/` to toggle ongoing AML monitoring for up to 1,000 profiles at a time, with dry-run and async batch support.

NEW INTEGRATION Adds `POST /v3/kyb/shareholders/` and `POST /v3/kyb/ubo/` for three-tier KYB registry data (Lite, Shareholders, UBOs) with per-response tier confirmation.

Adds `GET /v3/users/{vendor_data}/generate-pdf/` to return a branded PDF of a user's 20 most recent completed sessions.

Adds `sandbox_scenario` on `POST /v3/aml/` to simulate AML hits in sandbox without an end-user session.

Exposes `workflow_terminal_node_id` on the decision API via `include=events` to show which node ended a session.

+7 more

NEW INTEGRATION Digital-ID wallet integrations: MitID, Finnish Trust Network, Smart-ID, Mobile-ID, BankID Sweden

NEW DETECTION Document liveness scoring (portrait replacement, printed copy, screen replay) on every decision API result

NEW DETECTION Emulated/virtualized device detection for liveness captures

NEW DETECTION Workflow branch contradiction validation flags routing errors before publish

NEW INTEGRATION Didit MCP tools for session explanation and webhook management

↳ BREAKING ON UPGRADE

! `X-Signature-Simple` is deprecated — verify webhooks with `X-Signature-V2` instead.

WAS THIS USEFUL?



All the latest releases from the 1Password team

1Password shipped updates to its Environments feature on Linux and a round of sign-in prompt and diagnostics improvements in the browser extension.

↳ WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

01 Environments improvements in 1Password for Linux IMPROVED

68

Raises the 1Password Environments limit to **25** local **.env** file destinations per computer, up from the previous cap, and enables data export for users signed in via Unlock with SSO when an administrator permits it. After adding an environment variable or importing a **.env** file, users are returned automatically to the Environments overview screen, and the Delete button is now always available when removing a workload, showing an error if the entered name does not match the workload name.

– Names limit, SSO condition, and exact delete-confirmation behavior.

1Password for Linux 8.12.42 (2026-10-20)

02 Sign-in prompt upgrades in 1Password browser extension IMPROVED

62

The sign-in prompt now includes a search field to filter relevant Login items, automatically offers passkeys instead of passwords when a Login item has a passkey saved, supports Left and Right Arrow key navigation between a Login item's sign-in methods, and shows a single sign-on reminder even when the browser extension is not connected to the 1Password app.

– Lists four concrete sign-in prompt behaviors but only UI-level detail.

1Password in the browser 8.12.42 (2026-...

THINNER COVERAGE BELOW

03 Diagnostics report creator for Firefox NEW

26

Adds a diagnostics report creator for 1Password in Firefox.

– Single bare-name bullet with no mechanism or access path given.

1Password in the browser 8.12.42 (2026-...

WAS THIS USEFUL?



Nudge Security

Feedback



Discover AI agents built on Notion

NOTES

RANK

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

Nudge Security shipped native ServiceNow and Jira Cloud ticketing integrations with auto-close on finding resolution, added an Okta integration for SSO enrollment monitoring, and extended AI agent discovery to agents built on Notion.

WHAT SHIPPED • 4 FEATURES most completely described first ? what's the number?

01 Okta SSO enrollment analysis integration NEW 95

Adds an Okta integration via a read-only or application administrator API token, requiring access to `/api/v1/apps`, `/api/v1/apps/{application_id}/groups`, `/api/v1/apps/{application_id}/users`, `/api/v1/org/contacts/technical`, `/api/v1/users/`, `/api/v1/groups`, and `/api/v1/groups/{group_id}/users` endpoints. Connecting an Okta tenant enables monitoring of SSO enrollment status across the full SaaS estate, and the `Things to do > Onboard apps to Okta` playbook can automatically initiate SSO enrollment workflows.

...

Connect your Okta tenant to Nudge Security to start tracking SSO enrollment status across your SaaS estate.

1. In Okta admin console, go to Security > API > Tokens and click `Create token` using a read-only administrator or application administrator account. 2. Copy the generated token. 3. In Nudge Security, navigate to Things to do > Onboard apps to Okta and start or continue the playbook. 4. Enter your Okta API token and subdomain, then click `Test connection`. 5. Confirm status shows `Verified` to complete the integration.

– Lists exact API endpoints and a runnable playbook path [product docs](#)

02 ServiceNow ticketing integration NEW 90

Adds a ServiceNow integration configurable at `Settings > Ticketing > Connect ServiceNow`, usable alongside an existing Jira connection with per-ticket destination selection. Default field mapping under `Settings > Ticketing > ServiceNow` pre-populates ticket fields (Table, Urgency, Assigned to, Assignment group, Caller, Short description, Description) when creating tickets from findings. When a finding is resolved in Nudge, the linked ServiceNow ticket auto-closes to the mapped state with a work

note containing the finding name, resolution date, and a link back to Nudge; if the finding re-emerges, a new ticket is created while the original remains closed as a historical record.

Connect Nudge Security to ServiceNow and configure default field mappings so tickets created from findings are pre-populated correctly.

📌 In the Nudge Security console, go to Settings > Ticketing, select **Connect ServiceNow**, enter your ServiceNow instance URL, and authorize access. Then go to Settings > Ticketing > ServiceNow to configure your default field mappings (Table, Urgency, Assigned to, Assignment group, Caller).

Create a ServiceNow ticket from an open finding and have it auto-close when the finding is resolved, with full Timeline history in Nudge.

📌 Open a finding in Nudge Security, go to the Fix tab, select **Create ticket**, choose **ServiceNow**, review the pre-filled Short description and Description, adjust Table/Urgency/Assigned to as needed, then select **Create**. The ticket (e.g., INC0012345) is linked back to the finding and will automatically move to your mapped closed state when the finding resolves.

– Full config paths and field names, but UI steps not API [product docs](#)

03 Jira Cloud ticketing integration **NEW**

90

Adds Jira Cloud integration configurable at **Settings > Ticketing**, with a **Create Jira ticket** action on the **Fix** tab of any finding that pre-populates Summary, Description, Project, Work Type, Priority, Assignee, and Reporter from default field mappings set at **Settings > Ticketing > Jira**. The Jira ticket key and URL are stored on the finding and a **Jira ticket created** event with an inline link is logged in the finding Timeline; if a resolved finding re-emerges, a new Jira ticket is created while the original closed ticket is preserved as a historical record.

:

Connect Jira Cloud and configure default field mappings so tickets created from findings are pre-populated correctly.

📌 In Nudge Security, go to Settings > Ticketing, select **Connect Jira**, and authorize access to your Jira Cloud site. Then go to Settings > Ticketing > Jira to configure default field mappings (Project, Work Type, Priority, Assignee, Reporter).

Create a tracked Jira ticket from an open finding so engineers can work it in Jira while Nudge auto-closes the ticket on resolution.

📌 Open a finding > Fix tab > select **Create Jira ticket** > review and adjust pre-filled fields > select **Create**. The ticket key and URL are stored on the finding, and the

Timeline logs a **Jira ticket created** event with a deep link.

– Names config paths and fields, UI-based setup [product docs](#)

04 **AI agent discovery for Notion** NEW

60

Enables the **Notion AI agents** connected app, surfaced under the Identities > AI Agents tab, to automatically discover and inventory AI agents your team builds on Notion, including creator, components, and risk insights.

Enable Notion AI agent discovery so every agent your team builds surfaces automatically in the AI agents inventory with creator, components, and risk insights.

 In Nudge Security, go to Identities > AI Agents tab, then enable the **Notion AI agents** connected app to start ingesting agents into your inventory.

– Names the connected app and UI path but little mechanism detail

[Discover AI agents built on Notion](#)

WAS THIS USEFUL?



RANK

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

JumpCloud's biggest additions this window are Agent Uninstall Protection with maintenance tokens, MFA Frequency Controls for Conditional Access Policies, and new Custom Risk Policies and Patch Management Dashboards in the Admin Portal, alongside REST APIs for AI Agent management, an MCP Server for the User Portal, App Catalog update controls, and new Apple MDM app-privacy policies.

↳ WHAT SHIPPED • 12 FEATURES most completely described first ? what's the number?

01 Agent Uninstall Protection for endpoint agents NEW

85

JumpCloud now enables Agent Uninstall Protection by default on all devices running agent v2.179.2 or later, requiring a per-device maintenance token to uninstall the agent via standard methods on macOS, Windows, and Linux. Admins can exempt specific devices from this requirement using the **Disable Agent Uninstall Protection** policy for macOS, Windows, and Linux, for maintenance, troubleshooting, or imaging purposes.

Retrieve a maintenance token for a specific device before manually uninstalling the JumpCloud agent, so the removal is authorized without network connectivity.

📌 In the Admin Portal, go to Device Management > Devices, open the target device's details, and copy the maintenance token from Agent Uninstall Protection. Present the token when prompted during the uninstall.

– Names the policy, agent version, platforms and token workflow.

Release Notes 2026 (2026-10-16)

02 MCP Server for Users in JumpCloud AI NEW

80

A new **MCP Server for Users** toggle under **Settings > JumpCloud AI** lets users connect AI clients to User Portal capabilities via OAuth sessions, with admin controls to review and revoke active sessions and optionally configure an automatic session expiration window.

Enable the MCP Server for Users so staff can perform User Portal tasks directly from AI clients without leaving their tools.

📌 In the Admin Portal, go to Settings > JumpCloud AI and toggle on MCP Server for Users. Optionally configure an automatic session expiration window and review

active OAuth client sessions.

– Exact config path and steps given via usage example. Release Notes 2026 (2026-10-16)

03 Patch Management Dashboards in Admin Portal NEW 75

New Patch Management Dashboards under **Device Management > Patch Management** provide unified Overview, Windows, and Apple tabs, including KB-level CVE details, severity filtering, bulk remediation, and export to CSV or JSON.

– Names exact nav path, tabs and export formats. Release Notes 2026 (2026-10-16)

04 App Settings Privacy policies for Apple MDM NEW 75

New App Settings Privacy policies let admins manage app permissions on Apple devices: **User Level macOS App Settings Privacy** (macOS 27+, supervised Macs, assignable to users or user groups) and **Device Level iOS and iPadOS App Settings Privacy** (iOS/iPadOS 27+, supervised devices, assignable to devices or device groups). Both support permissions including Accessibility, Bluetooth, Camera, Dictation, Local Network, Location, and Microphone.

– Names both policies, OS versions and permission list. Release Notes 2026 (2026-10-16)

05 Custom Messaging in User Portal NEW 70

A new **Custom Messaging** feature under **Settings > User Portal > Custom Messaging** lets admins create login dialogs and announcement banners in the User Portal, targetable by user group with configurable severity.

– Exact config path and targeting options named. Release Notes 2026 (2026-10-16)

06 Asset visibility and acknowledgement settings for end users NEW 70

New **Asset Visibility** and **Asset Acknowledgement** settings under **Device Management > Asset Management > Settings and Fields > General Settings** give end users visibility into assigned assets and record their acknowledgements in Asset History.

– Exact nested config path given for both settings. Release Notes 2026 (2026-10-16)

07 Update notification controls in App Catalog IMPROVED

70

Admins can toggle **Disable end user update notifications** when adding or editing an App Catalog application, applying to Auto and Self Service delivery (existing apps keep notifications enabled, and Auto-Update still forces updates after seven days). Users can also start available updates for Required applications directly from the Updates section in the Self Service App Catalog, regardless of the Auto-Update setting.

– Names the exact setting, delivery modes and default behavior. [product docs](#)

08 Custom Risk Policies for logins NEW

65

Admins can create Custom Risk Policies for Admin Portal, User Portal, device, and SSO application logins — selecting risk factors per user group or SSO application, adjusting their score contributions, and setting policy precedence.

– Explains scope and configuration options but no exact UI path.

[Release Notes 2026 \(2026-10-16\)](#)

09 REST APIs for managing AI Agents and Agent Groups NEW

60

New REST APIs let admins programmatically register, retrieve, update, suspend, activate, and delete AI agent identities, credentials, groups, memberships, and application assignments.

– Lists API actions but no endpoint paths or methods. [Release Notes 2026 \(2026-10-16\)](#)

10 Folders for organizing apps in User Portal NEW

60

End users can now organize assigned applications into private personal folders on the **All Applications** page via the new **Folders** tab.

– Exact page and tab named for navigation. [Release Notes 2026 \(2026-10-16\)](#)

THINNER COVERAGE BELOW

11 Custom Banned Password Lists NEW

45

Organization administrators can now create a Custom Banned Password List to block specific words, terms, and phrases from user passwords.

– No config path or scope detail beyond the feature name. [Release Notes 2026 \(2026-10-16\)](#)

12 MFA frequency controls in Conditional Access Policies NEW

45

Conditional Access Policies now support MFA Frequency Controls, letting admins define how often users must re-authenticate with MFA.

– No config path, default interval, or configuration steps given. [product docs](#)

↳ ALSO FROM THESE RELEASES

- .
- .
- .
- .

WAS THIS USEFUL?



Duality Technologies

Feedback



Is Anyone on This Flight a Threat? What If You Could Find Out Without Revealing a Single

BLOG ↗

RANK

Duality technologies is a healthcare, finance, and government organizations use Duality to analyze sensitive data across silos - without ever exposing it. See how it works.

Duality launched Secure Query, letting agencies screen passenger manifests against multiple classified watchlists using encrypted queries without exposing passenger identities or watchlist contents, with potential extension to border screening, sanctions checks, and law-enforcement collaboration.

↳ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 Duality Secure Query for watchlist screening NEW 53

Duality Secure Query lets authorized organizations screen passenger manifests against multiple intelligence-agency watchlists using encrypted queries, without exposing passenger names to the queried agencies or sharing the underlying watchlist data. Configured screening workflows also support source-attribution concealment, returning an authorized match alert without revealing which participating agency holds the relevant intelligence.

- Names the product and mechanism but no API, config or limits given

Is Anyone on This Flight a Threat? What...

02 Encrypted query use cases beyond aviation NEW 22

Encrypted query collaboration is described as extending beyond aviation screening to border screening, financial sanctions checks, cross-border law-enforcement investigation, and allied-defense threat identification across national and classification boundaries.

- Only prose list of potential domains, no concrete mechanism given

Is Anyone on This Flight a Threat? What...

WAS THIS USEFUL?



Cato Networks

Feedback



Stop Guessing. Find the Root Cause with Cato DEM. (2026-10-20)

BLOG

RANK

Cato networks is a cato SASE Cloud converges SD-WAN, security, global backbone, and remote access into a global cloud-native service. WE ARE SASE™

Cato Networks added Root Cause Analysis to Cato DEM, pulling together endpoint, network adapter and path-level evidence to help pinpoint why digital experience is degraded.

WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Root Cause Analysis in Cato DEM

NEW

64

Cato DEM gains a Root Cause Analysis (RCA) workflow that correlates endpoint, network, application, and historical path evidence in a single investigation to identify the likely fault domain. It includes Network Adapter Metrics (adapter type, BSSID, channel, noise, and signal strength) to assess local Wi-Fi conditions, and Client Traceroute, which provides historical hop-by-hop path evidence capturing latency, loss, and path behavior, especially for intermittent and remote-user issues. RCA scope extends beyond a single user to correlate experience data across multiple users to reveal whether an issue is isolated or shared and its blast radius, and works across both Cato SASE and supported third-party network environments without requiring replacement of existing infrastructure.

– Describes mechanism and named metrics but no UI path or command

Stop Guessing. Find the Root Cause with...

WAS THIS USEFUL?



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

Tailscale introduced MagicRoute (beta), a new automatic path-selection system for redundant network routes, and replaced the admin console's Regional Routing toggle with a unified Route selection configuration surface.

←→ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 MagicRoute automatic redundant-path selection NEW

62

MagicRoute (beta) automatically chooses between redundant subnet routers, app connectors, exit nodes, and Tailscale Service hosts to provide resilient routing across a tailnet. It can be enabled in the admin console under **Settings > Route selection**.

Enable MagicRoute for your tailnet to let Tailscale automatically pick the best available subnet router or exit node when multiple redundant paths exist.

📌 In the admin console, go to Settings > Route selection and enable MagicRoute (beta).

– Names mechanism and components but no limits or internals

MagicRoute (2026-10-19)

THINNER COVERAGE BELOW

02 Route selection replaces Regional Routing toggle

BREAKING

48

The admin console's Regional Routing toggle has been removed and replaced by a new **Route selection** configuration surface governing tailnet-wide route policy. Existing setups relying on the old toggle must be reconfigured under this new surface.

– Clear migration note but no further mechanism detail

MagicRoute (2026-10-19)

←→ BREAKING ON UPGRADE

! The Regional Routing toggle in the admin console is removed and replaced by the **Route selection** configuration; existing setups relying on that toggle must be reconfigured under the new **Route selection** surface.

WAS THIS USEFUL?



AGENT

Runlayer

Feedback



1.29.0 - iMessage agents and recorder detection (2026-10-20)

BLOG ↗

LATEST

STABLE

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

Adds iMessage agent channel, GitHub App integration for agents, ExtendDB NetworkPolicy isolation, RunlayerInstance CR image overrides, and semantic customer versioning.

Adds iMessage channel for agents: users can text an agent from Messages on iPhone, iPad, or Mac after linking their number to their Runlayer account; admins choose which agent answers the linked number.

Adds GitHub App integration for agents: workspace admins can connect a GitHub App and grant agents access to selected repositories via short-lived tokens.

Adds `RunlayerInstance` CR overrides for ExtendDB init/reset Jobs' helper container images (kubectl, Postgres client, AWS CLI) on self-hosted Kubernetes, removing the need to wait on an operator release.

Adds ExtendDB ingress `NetworkPolicy` for self-hosted deployments that scopes ExtendDB to only the linked MCP namespace.

Adds Postgres connection-budget sizing guidance and an image-bump/recovery runbook for self-hosted Kubernetes operators enabling ExtendDB.

+2 more

NEW DETECTION

AI meeting recorder and ambient notetaker detection in AI Watch Detect

NEW DETECTION

Computer-use desktop agent detection

NEW DETECTION

pip- and npm-installed computer-use tool detection

UPDATED DETECTION

AI Watch Detect Capabilities chips for installed AI clients

NEW DETECTION

Gemini desktop app inventoried as AI Watch client

NEW INTEGRATION

Microsoft 365 Work IQ connector

➡ BREAKING ON UPGRADE

! The separate `Goal achieved` tile has been removed; goal outcomes and failure reasons are now only available in Agent Insights.

1.28.1575 - Computer-use detection and session stop (2026-10-20)

[BLOG ↗](#)

STABLE

Adds session stop controls (Shift+K), scoped agent token enforcement (`FF_DISABLE` / `FF_ENABLE`), credential revocation on run end, and redaction-aware stream replay.

Security, IT, and Super Admins can stop a live agent session and resume it from the Sessions page session drawer, or with `Shift+K`, with a required reason and full audit trail.

Scoped agent or on-behalf-of tokens (requested with `resource`) now return `403 invalid_target` on out-of-scope servers for all agent accounts; self-hosted operators control this with `FF_DISABLE=agent-token-audience-enforcement` or `FF_ENABLE=agent-token-audience-enforcement`.

Agent run credentials are revoked the moment a run finishes rather than remaining valid until natural expiry.

Agent run live-stream replay enforces workspace redaction rules; viewers must hold the `unredact` permission and confirm before content is replayed.

Agents offer an inline connector card in chat when a needed app is missing, eliminating a settings-page detour.

+1 more

↳ BREAKING ON UPGRADE

- ! Scoped agent or on-behalf-of tokens are now rejected with `403 invalid_target` on servers outside their declared `resource` scope for ALL agent accounts — self-hosted operators can revert with `FF_DISABLE=agent-token-audience-enforcement`.
- ! Exchanging a scoped agent token for an on-behalf-of token now requires a `resource` within the original token's scope; requests without one return `400 invalid_target`.
- ! The MDM configuration card's `Last connected` label and status badge are removed.
- ! The `Turn on Shadow AI detection` onboarding step has been removed.

WAS THIS USEFUL?



NVIDIA OpenShell

[Feedback](#)

v0.1.5

[NOTES](#)[RANK](#)

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

OpenShell v0.1.5 cuts policy-change propagation to ~100 ms with push-based config delivery, adds OCSF-over-OTLP export and virtiofs host mounts for MicroVM sandboxes, and ships a cluster of breaking changes spanning the Kubernetes driver, Go SDK, snap packaging, and OCSF log file naming.

↳ WHAT SHIPPED • 13 FEATURES most completely described first ? what's the number?

01 Kubernetes pod annotation and label allowlisting

BREAKING

88

Adds `allowed_pod_annotations` and `allowed_pod_labels` settings in `[openshell.drivers.kubernetes]` to allowlist pod annotation and label keys used in sandbox templates. Breaking: sandbox templates setting pod annotations or labels are now rejected unless the operator allowlists the keys; API and SDK callers setting `SandboxTemplate.labels` or `SandboxTemplate.annotations` (including recovery of timed-out sandboxes) are affected, while the `openshell` CLI and `--label` metadata labels are not.

– Named settings, scope of impact, and migration path v0.1.5

02 Go SDK sandbox creation API changed to functional options

83

BREAKING

The Go SDK's `SandboxInterface.Create()` and `CreateFromTemplate()` now take functional options (`WithLabels()`, `WithAnnotations()`, `WithServiceExposures()`) instead of a positional labels argument; the `CreateOptions` struct has been removed, and call sites must be updated.

– Exact API signatures and required migration named v0.1.5

03 New sandbox create flags for annotations and env passthrough

NEW

81

Adds `--annotation KEY=VALUE` repeatable flag to `openshell sandbox create` for setting sandbox annotations from the CLI, and `--env-from KEY[=ENVVAR]` flag to pass sandbox environment variables from the caller's environment without putting values on the command line.

Pass secrets from the local environment into a new sandbox without exposing them in shell history or process listings.

```
$ openshell sandbox create --name agent-run --env-from  
AWS_ACCESS_KEY_ID --env-from AWS_SECRET_ACCESS_KEY --annotation  
team=security -- python agent.py
```

– Exact flags and a runnable command shown v0.1.5

04 Push-based config delivery NEW

80

Adds `config_delivery_mode = "push"` supervisor setting to receive configuration changes pushed over the gateway session in ~100 ms instead of up to 10 seconds; polling remains the default.

Reduce policy-change propagation from up to 10 s to ~100 ms by enabling push-based config delivery on the supervisor.

```
toml  
  
config_delivery_mode = "push"
```

– Named config key with clear before/after numbers v0.1.5

05 MicroVM virtiofs host bind mounts NEW

80

Adds `enable_bind_mounts = true` setting in `[openshell.drivers.vm]` to allow MicroVM sandboxes to mount host directories through virtiofs; supports up to 3 mounts on x86_64 or 32 on aarch64, with QEMU-backed and GPU launches not supported.

– Config key with concrete mount limits and caveats v0.1.5

06 Snap user-gateway service no longer enabled globally

BREAKING

80

The snap no longer enables the user gateway service globally; after a direct snap install, or if an older revision enabled it globally, operators must run `sudo systemctl disable --global snap.openshell.user-gateway.service` once, then `systemctl --user enable --now snap.openshell.user-gateway.service` from the single account that should run the gateway.

– Exact migration commands make this directly runnable v0.1.5

07 OCSF events exported over OTLP NEW

78

Adds `[openshell.gateway.otlp.ocsf]` config block with `enabled = true` to export OCSF events as OTLP log records, allowing traces and OCSF events to be sent to

separate OTLP endpoints; also adds `OPENSHELL_OCSF_JSON=1` gateway audit sink, now available on all platforms.

Stream OCSF audit events into your OTLP pipeline (e.g. a separate security observability backend) without mixing them with traces.

```
toml

[openshell.gateway.otlp.ocsf]
enabled = true
endpoint = "https://ocsf-collector.example.com:4317"
```

– Config block and env var named, endpoint routing explained v0.1.5

08 OCSF log files renamed to .jsonl

BREAKING

75

Supervisor OCSF log files are renamed from `openshell-ocsf.YYYY-MM-DD.log` to `openshell-ocsf.YYYY-MM-DD.jsonl`; collector glob patterns must be updated from `openshell-ocsf.*.log` to `openshell-ocsf.*.jsonl`, and old `.log` files are not migrated.

– Exact file name patterns and migration note given v0.1.5

09 Compute driver bootstrap descriptor field renamed

BREAKING

65

External compute drivers that provision `openshell-sandbox` must update their bootstrap descriptor field from `outer_fence` to `provisioning_evidence`.

– Exact field rename named, limited context otherwise v0.1.5

THINNER COVERAGE BELOW

10 Configurable gateway data volume sizing in Helm

NEW

55

Adds `server.dataStorageSize` and `server.dataStorageClass` Helm values to configure the gateway data PVC.

– Named Helm values but no further detail v0.1.5

11 Reserved annotation prefix rejected on sandbox APIs

BREAKING

55

The gateway now rejects annotation keys with the reserved `internal.openshell.ai/` prefix on `CreateSandbox` and `UpdateConfig`.

– Named endpoints and prefix but thin explanation v0.1.5

12 Per-workspace sandbox mutation locks IMPROVED

50

Gateway mutation locks are now scoped per workspace and sandbox with bounded waits, preventing bursts of supervisor reconnects from queuing behind a single fleet-wide lock.

– Describes mechanism but no flag or config to invoke v0.1.5

13 Docker supervisor bridge networking mode NEW

45

Adds `supervisor_network_mode` setting for the Docker supervisor to enable bridge networking; host networking remains the default.

– Named setting but no detail on values or behaviour v0.1.5

↳ BREAKING ON UPGRADE

- ! Supervisor OCSF log files are renamed from `openshell-ocsf.YYYY-MM-DD.log` to `openshell-ocsf.YYYY-MM-DD.jsonl`; update any collector glob patterns from `openshell-ocsf.*.log` to `openshell-ocsf.*.jsonl`. Old `.log` files are not migrated.
- ! On Kubernetes, sandbox templates that set pod annotations or labels are now rejected unless an operator allowlists the keys with `allowed_pod_annotations` and `allowed_pod_labels` in `[openshell.drivers.kubernetes]`. API and SDK callers setting `SandboxTemplate.labels` or `SandboxTemplate.annotations` (including recovery of timed-out sandboxes) are affected; the `openshell` CLI and `--label` metadata labels are not.
- ! The gateway now rejects annotation keys with the reserved `internal.openshell.ai/` prefix on `CreateSandbox` and `UpdateConfig`.
- ! The Go SDK's `SandboxInterface.Create()` and `CreateFromTemplate()` now take functional options (`WithLabels()`, `WithAnnotations()`, `WithServiceExposures()`) instead of a positional labels argument; the `CreateOptions` struct has been removed. Call sites must be updated.
- ! External compute drivers that provision `openshell-sandbox` must update their bootstrap descriptor field from `outer_fence` to `provisioning_evidence`.
- ! The snap no longer enables the user gateway service globally; after a direct snap install or if an older revision enabled it globally, run `sudo systemctl disable --global snap.openshell.user-gateway.service` once, then `systemctl --user enable --now snap.openshell.user-gateway.service` from the single account that should run the gateway.

WAS THIS USEFUL?



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

h5i v0.5.1 adds a full browser-automation and policy-checking proxy workflow, three installable plugins for websec, recon, and automated test replay, plus a new agent skill interface and GPU-backed app engine.

↳ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 Websec plugin for HTTP capture, replay, and blind-oracle testing

NEW

93

`h5i plugin install websec --from <path>` installs the `h5i-websec` workbench, enabling HTTP capture, modification, replay, and comparison for vulnerability investigation. When a websec replay goes out unauthenticated, the tool now prompts the user toward supplying `--as` to replay under a specific identity, and a new blind-oracle workflow is documented under `skill/websec`.

Install all three optional plugins (websec, recon, test) in one shot using the convenience install script.

```
$ curl -fsSL https://h5i.dev/install.sh | sh -s -- --websec --recon --test
```

– Install command, flag, and workflow all named with before/after behaviour v0.5.1

02 Test plugin for attack-flow replay with coverage gating NEW

90

`h5i plugin install test --from <path>` installs `h5i-test`; `h5i test <dir> --target <url>` replays repository-owned attack flows, with optional `--min-coverage` for OpenAPI coverage enforcement, and writes JSON and JUnit output.

Run security regression tests against a local app in CI, enforcing a minimum OpenAPI coverage threshold so gaps fail the build.

```
$ h5i test .h5i-tests/tests --target http://localhost:3000 --min-coverage 80
```

Install all three optional plugins (websec, recon, test) in one shot using the convenience install script.

```
$ curl -fsSL https://h5i.dev/install.sh | sh -s -- --websec --recon --test
```

– Full command, flag, and output formats with a runnable CI example v0.5.1

03 Browser automation and session monitoring NEW 83

New `h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` commands drive a Chromium session through a built-in policy-checking proxy that logs every request before it moves. The proxy now decodes zstd-encoded responses and stores proxy-observed bodies decoded, enabling replay of modern compressed traffic. A new `h5i ui` command opens a local dashboard for monitoring browser and session activity.

Capture and review all HTTP requests made during a browser session against a target, to identify attack surface before replaying.

```
$ h5i browser open https://example.com
h5i browser requests
```

– Names four commands and proxy behaviour, with a runnable example v0.5.1

04 Recon plugin for endpoint discovery NEW 60

`h5i plugin install recon --from <path>` installs `h5i-recon`, which records endpoints, parameters, and candidate URLs discovered through crawling and JavaScript analysis.

Install all three optional plugins (websec, recon, test) in one shot using the convenience install script.

```
$ curl -fsSL https://h5i.dev/install.sh | sh -s -- --websec --recon --test
```

– Names install command and plugin but little mechanism beyond that v0.5.1

THINNER COVERAGE BELOW

05 Skill install subcommand for agent runtime NEW 45

`h5i skill install` writes the agent-facing skill interface where the agent runtime expects it.

– Command named but no detail on mechanism or output v0.5.1

Adds GPU backend support to the app engine with generalized inputs and batch output validation.

– Described only in prose, no flag or command to act on `v0.5.1`

WAS THIS USEFUL?



Docker Sandboxes

[Feedback](#)[v0.49.1](#)[NOTES](#)[RANK](#)

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

Docker Sandboxes v0.49.1 adds offline kit validation and build-arg support, new daemon log/browser/proxy settings, broader proxy-aware networking for cloud and MCP commands, safer secret and policy commands, and richer CLI diagnostics.

WHAT SHIPPED · 10 FEATURES most completely described first ? what's the number?

01 Offline validation and build-arg support in kit workflow

IMPROVED

90

`sbx kit validate` and `sbx kit inspect` now accept `--offline` to check a local directory or ZIP kit without starting the daemon or signing in. Sandboxes created from local or Git v3 kits now honor build-time `--kit-arg` values, including scoped overrides. `sbx kit sign` warns when Git would change the signed file contents and invalidate signatures for `git+` references, and `sbx settings set` rejects incorrectly structured values for `kit.allowedSources` and `kit.trustedSigners` before saving them.

Validate a local kit archive before deploying, without starting the daemon or authenticating — useful in CI.

```
$ sbx kit validate --offline ./my-kit.zip
```

— Names every flag and config key with a runnable offline example. [v0.49.1](#)

02 New daemon log, browser, and proxy settings

NEW

85

`sbx settings set` gains `daemon.log.maxSizeMB`, `daemon.log.maxBackups`, and `daemon.log.maxAgeDays` to control daemon log size and retention, `browser.enabled` to control whether sandboxes can open the host browser, and `proxy.sshDirect` (set to `true`) to connect cloud SSH directly on networks that permit direct SSH connections.

Cap daemon logs to 50 MB per file, keep 5 backups, and expire files older than 30 days to prevent disk fill on long-running hosts.

```
$ sbx settings set daemon.log.maxSizeMB=50
daemon.log.maxBackups=5 daemon.log.maxAgeDays=30
```

Force cloud SSH through a direct connection on a corporate network that allows outbound SSH, bypassing the HTTP proxy.

```
$ sbx settings set proxy.sshDirect=true
```

– All keys named with concrete runnable set commands. v0.49.1

03 Safeguards for secret removal and overwrite IMPROVED

75

`sbx secret rm` now returns an error when no secret matches `--registry`, `--placeholder`, `--host`, or `--env-name`, and `sbx secret set-custom` warns before replacing a different host or environment-variable mapping under the same placeholder.

– Names all flags affected, no runnable example provided. v0.49.1

04 Proxy-aware OAuth discovery and cloud commands IMPROVED

70

`sbx mcp add` now completes OAuth discovery on networks that require an HTTP proxy without needing `--skip-ssrf-check`, and cloud commands, `sbx move`, the TUI, and cloud SSH now honor configured proxies, including system and PAC proxies.

– Names commands and proxy types but lacks a usage example. v0.49.1

05 Expanded diagnostics in `sbx diagnose` IMPROVED

65

`sbx diagnose` now reports running CLI and shim processes, warns when global network policy is not initialized, and detects a ZFS setting that can make image builds unusually slow.

– Lists concrete checks but no command example. v0.49.1

06 `-q` flag for policy and settings list commands NEW

60

`sbx policy ls`, `sbx settings ls`, and `sbx policy approval ls` now accept `-q` to print only rule IDs, setting keys, or approval IDs.

– Flag and commands named but no example given. v0.49.1

07 Detailed error locations in `sbx env` IMPROVED

60

`sbx env` reports unknown keys and invalid values with their source file, line, column, and field.

– Explains mechanism of error reporting, no example shown. v0.49.1

08 **`sbx policy refresh` applies immediately** IMPROVED 55

`sbx policy refresh` immediately refreshes policy instead of waiting for the next polling interval.

– Clear before/after behaviour but no further mechanism. `v0.49.1`

09 **C/C++ toolchain in full templates** NEW 40

Full templates now include a C/C++ toolchain for compiling native code inside the sandbox.

– States the addition with no further detail. `v0.49.1`

10 **Linux tarball installer works under AppArmor** IMPROVED 40

The Linux tarball installer now works for regular users on systems with AppArmor.

– Bare fix description with no mechanism. `v0.49.1`

WAS THIS USEFUL?  

◆ AI Red Teaming & Testing

Giskard Feedback ? giskard-core/v1.0.3 NOTES RANK

Giskard tests and evaluates LLM applications for performance, bias, and security issues.

Giskard Core v1.0.3 introduces an opt-in strict template policy for agent prompt enforcement.

←→ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 **Strict template policy for agents** NEW 45

Adds an opt-in strict template policy to `giskard.agents`, controlling how prompt templates are enforced during agent runs.

– Names the module and setting but no mechanism detail `giskard-core/v1.0.3`

WAS THIS USEFUL?  

OTHER / UNCATEGORIZED

ArmorCode

[Feedback](#)

announced · Oct 21, 2026

[BLOG](#)[RANK](#)

ArmorCode is the independent control plane where AI agents help find, prioritize, and fix vulnerabilities across apps, cloud, infra, and AI.

ArmorCode's biggest move this window is the launch of ArmorCode Zero, an agentic layer that triages findings, groups them into unified fixes, drafts pull requests, and rescans to verify remediation — alongside a new Context Risk Graph for prioritization and a conversational chart builder.

↳ **WHAT SHIPPED · 4 FEATURES** most completely described first [? what's the number?](#)

01 **Agentic triage and remediation via ArmorCode Zero** **NEW**

64

ArmorCode Zero reads every finding, writes a triage recommendation, and groups findings by the single change that resolves them, then drafts pull requests and rescans to verify results. On one tenant this reduced 28,600 findings to roughly 8,900 actionable fixes.

– Explains mechanism and gives concrete before/after numbers, but no exact UI or API path.

[launch-20261021-d9b49697](#)

THINNER COVERAGE BELOW

02 **ArmorCode Context Risk Graph** **NEW**

47

A shared knowledge layer that connects findings to assets, code, owners, business context, services, and controls to drive prioritization inference across the platform.

– Names the entities it connects but gives no mechanism for how inference works.

[launch-20261021-d9b49697](#)

03 **Conversational chart-building interface** **NEW**

38

Security teams can describe a visualization in plain language within Zero and it generates the corresponding chart.

– Describes the interaction pattern but no concrete chart types or limits.

[launch-20261021-d9b49697](#)

04 **Secrets redaction and bring-your-own model keys** **NEW**

37

Secrets are redacted before any model call is made, and teams can supply their own model API keys rather than relying on a shared one.

– States the safeguard and option but no config key or setup steps.

launch-20261021-d9b49697

WAS THIS USEFUL?



// END OF TRANSMISSION ■

The Cyber Toolchain · curated daily · October 21, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe**

Feedback on this issue