THE CYBER TOOLCHAIN

NO. 138

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 13, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 172 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 59 tools

▾ // THE BRIEF

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

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

Three unrelated things stand out. GitGuardian SaaS planted a Kubeconfig honeytoken that fires when a secret is lifted off a developer

vim ? help

laptop rather than a repo; Nudge Security now catches a breached password at the moment of login via its browser extension; and Prowler 5.49.0 matches emerging-threat signatures against running AWS, Azure and GCP compute instances.

DETECT

Catch credential theft at the machine, not in the repo

A Kubeconfig honeytoken is a fake cluster credential that nobody should ever use, so any authentication attempt against it is a high-confidence signal that something read a developer's home directory — covering the leak path that repo scanning cannot see. Nudge's extension checks credentials at login, so a reused-and-breached password surfaces as an account finding instead of waiting on a mandatory rotation cycle, and the Detection Engine's new provider detectors plus Bedrock long-term key and SSH host validation improvements cut down the triage on what it does flag.

GitGuardian SaaS · Nudge Security · GitGuardian Detection Engine

PROTECT

Hand out a credential for one request instead of baking it into the sandbox

OpenSandbox 1.2.1 brokers credentials per request and enforces egress scoped to specific ports, so untrusted code gets a short-lived secret and a narrow network path rather than an environment variable that persists for the life of the VM. Infisical's JWKS publishing lets relying parties verify KMS-signed tokens without a shared secret, and its CRL controls for internal CAs make revoking an issued cert actually enforceable; Passbolt adds passkey MFA and encrypted metadata, so even vault item names stop leaking to the server.

OpenSandbox · Passbolt

PROTECT

Know when a cloud instance matches a live campaign, and when an SBOM is lying to you

Prowler 5.49.0 adds emerging-threat signature matching for compute instances across the three major clouds with threat-based filtering in the Cloud API — so posture findings can be sorted by what is being exploited now rather than by generic severity. The cyclonedx-gradle-plugin's new build-safety flag fails generation when dependency resolution is

> incomplete, which stops the quietly truncated SBOM that makes a downstream vulnerability scan look clean.
>
> Prowler

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

xalgorix

Feedback ⓘ v4.6.188 NOTES ↗ RANK

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

> Xalgorix hardened SSRF protections across both its HTTP request tools and scan configuration fields, added automatic Go memory limit scaling tied to host RAM, and made resuming paused scans more predictable.

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

01 Automatic Go memory limit scaling with host RAM IMPROVED

95

The automatic Go soft memory limit now scales with host RAM: hosts up to 12 GB keep 45% of RAM (capped at 4 GB), while larger hosts get one-third of RAM with a 32 GB ceiling, overridden in that order by `XALGORIX_GO_MEM_LIMIT_MB` or `GOMEMLIMIT`. An empty `GOMEMLIMIT=` is now treated as unset rather than disabling the automatic limit.

Cap memory on a 32 GB host to 8 GB instead of the automatic one-third-of-RAM limit, overriding the new scaling formula.

```
$ XALGORIX_GO_MEM_LIMIT_MB=8192 xalgorix --web
```

— Exact thresholds, env vars, and a runnable override command v4.6.187

02 SSRF and credential-scoping validation for scan fields

85

BREAKING

A scan request's `discord_webhook` pointing at the local machine, dashboard listener, link-local, multicast, or cloud metadata address now gets a `400` (or a WebSocket `error` event) instead of proceeding; webhooks already saved in queue state or schedules are checked when sent and a refused one is logged. An `api_base` set to a host outside the stored credential's endpoint without a matching `api_key` also gets a `400`; supplying a per-request `api_key` replaces the stored key pool for that scan.

Send a scan notification to an internal relay on your LAN while confirming the webhook is accepted before the scan starts.

```
$ curl -s -X POST http://127.0.0.1:9137/api/scan \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"target": "https://staging.example.com",
"discord_webhook": "http://192.168.1.50:8080/relay"}'
```

— Names fields, status codes and an example request v4.6.187

03 Predictable resume handling for paused scans IMPROVED

70

Resuming a refused or still-shutting-down paused scan now behaves predictably: a refused resume answers `409` and leaves the scan resumable, and a scan resuming while still shutting down waits up to 3 seconds before returning `409` with a retry hint.

— Specific status codes and timing but no example call v4.6.187

04 SSRF guards on http_request and send_request tools NEW

60

The `http_request` and `send_request` tools now refuse requests whose URLs point at the local machine or the dashboard listener, matching the rule already applied to other gated tools. Each redirect target is checked against the same local-machine rule before being followed; when a redirect is refused, the tool returns the redirect response with an explanation rather than following it.

— Describes mechanism but no runnable flag or endpoint given v4.6.188

└──▷ BREAKING ON UPGRADE

! A scan request that sets a `discord_webhook` pointing at the local machine, dashboard listener, link-local, multicast, or cloud metadata address now gets a 400 (or a WebSocket `error` event) instead of proceeding; webhooks already saved in queue state or schedules are checked when sent and a refused one is logged.

! A scan request that sets `api_base` to a host outside the stored credential endpoint without supplying `api_key` now gets a 400 instead of proceeding.

! An empty `GOMEMLIMIT=` environment variable is now treated as unset and no longer disables the automatic memory limit.

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Bruce

Feedback ⓘ 1.17 NOTES ↗

Bruce is a predatory ESP32 firmware for M5Stack, Cardputer and LilyGo devices.

Bruce 1.17 adds WiFite WPS, SDR radio takeover, Wire Tools menu, compiled JS, UART JS API, DuckyScript DEFINE, and a new grid UI.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 1.17 https://github.com/BruceDevices/firmware.git
# already have the repo? check out this version:
$ git checkout 1.17
```

Add native WiFite (WPS) app to the Wi-Fi menu.

NEW SUPPORTED PLATFORM Add SDR raw I/Q takeover of the ESP32-S3 WiFi radio.

NEW INTEGRATION Add a Wire Tools menu with I2C serial commands, a continuity tester, logic analyzer, SUMP logic server, and I2C bus discovery tool.

Add compiled JavaScript support.

Add JavaScript UART API.

+10 more

NEW DETECTION 36 new SubGHz protocol decoders

NEW SUPPORTED PLATFORM ESP32-S3 SuperMini with ST7789 board support

NEW SUPPORTED PLATFORM Waveshare ESP32-C5 N16R8 board support

NEW SUPPORTED PLATFORM Waveshare ESP32-C5-Touch-LCD-3.5 board support

NEW SUPPORTED PLATFORM Lilka v2.3 board support

NEW SUPPORTED PLATFORM Freenove FNK0104B board support

NEW SUPPORTED PLATFORM Freenove FNK0103 board support

NEW SUPPORTED PLATFORM LilyGO TTGO T4 v1.3 board support

NEW SUPPORTED PLATFORM LILYGO T-Display-C5 board support

NEW SUPPORTED PLATFORM ESP32-S3 N16R8 ST7789V board support

NEW SUPPORTED PLATFORM M5Stack Mako and Pebble board support

NEW SUPPORTED PLATFORM CYD ESP32-C5 port with Dual-Band Analyzer

NEW SUPPORTED PLATFORM T-Watch Ultra support

NEW INTEGRATION ProtoPirate integrated into native RF

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

Datadog pathrunner

Feedback ⓘ v1.3.0 NOTES ↗ RANK

> pathrunner v1.3.0 adds identity export capabilities and a new interactive shell command.

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

01 **Identity export subcommand** NEW 75

A new `identity export` subcommand is available under the `identity` command tree in both the REPL and CLI, letting users export captured identities to hand credentials off to another tool or preserve them outside the workspace.

Export a captured identity after a successful exploit to hand credentials off to another tool or preserve them outside the workspace.

```
$ pathrunner> identity export
```

— Names exact subcommand and use case, runnable example given. v1.3.0

THINNER COVERAGE BELOW

02 **Interactive shell command** NEW 40

A new `shell` command starts an interactive shell session directly.

— Names the command but no further mechanism or example. v1.3.0

WAS THIS USEFUL? 👍 👎

◆ Wireless & RF

airgeddon

Feedback ⓘ v12.10 NOTES ↗ RANK

Airgeddon this is a multi-use bash script for Linux systems to audit wireless networks.

> airgeddon v12.10 adds conflict detection for wireless interfaces across concurrent multi-instance sessions and switches the multi-instance orchestrator file to a colon-delimited format.

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

01 **Colon-delimited multi-instance orchestrator file format**

40

BREAKING

The multi-instance orchestrator file now uses a colon-delimited format; existing orchestrator files from prior versions may need to be regenerated or updated.

— Names format change and migration note but no file path or example. v12.10

02 **Interface conflict detection in multi-instance mode** NEW

35

airgeddon now detects wireless interface conflicts when running in multi-instance mode, preventing overlapping interface use across concurrent sessions.

— Describes what it does but not detection mechanism. v12.10

└──▷ **BREAKING ON UPGRADE**

! The multi-instance orchestrator file format has changed to colon-delimited; existing orchestrator files from prior versions may need to be regenerated or updated.

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

SpecterOps SharpHound

Feedback ⓘ v2.18.0 NOTES ↗

Sharphound is a C# Data Collector for BloodHound

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.18.0
https://github.com/SpecterOps/SharpHound.git
# already have the repo? check out this version:
$ git checkout v2.18.0
```

NEW DETECTION DisabledExtensions CA registry value collection

NEW DETECTION RPC encryption CA registry key collection

WAS THIS USEFUL? 👍 👎

julius

Feedback ⓘ v1.7.3 NOTES ↗

└──▷ GET THIS VERSION ▾

Go

```
$ go install github.com/praetorian-inc/julius/cmd/julius@v1.7.3
```

git

```
$ git clone --branch v1.7.3 https://github.com/praetorian-
inc/julius.git
# already have the repo? check out this version:
$ git checkout v1.7.3
```

Go module path matches the repo · checked 13 Nov Wrong command?

UPDATED DETECTION MCP probe updated to 2026-07-28 spec

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ WAF, Bot & API Protection

Akto

Feedback ⓘ

Block AI Agents and Vendors by Team, Role, Device and User NOTES ↗

Akto discovers APIs, MCP servers and AI agents, tests and red-teams them, and enforces runtime guardrails and posture management.

Akto adds scoped AI agent blocking by team/role/device/user, host and path allowlisting, guardrail violation evidence and reviewer assignment, multi-turn attack remediation guidance, and on-demand Endpoint Shield Diagnostics.

Block unapproved AI agents org-wide or scope the block to specific teams, roles, devices, or users so individual groups keep the agents they rely on.

Allow or block browser-based AI agents and apps using the LLMs tab in Audit Data.

Permit traffic only to specified hosts and paths with host and path allowlisting — every request that does not match is blocked.

See the exact request and response evidence that triggered each guardrail violation, highlighted in the Guardrails dashboard.

Assign owners, add notes, and follow case timelines for every guardrail violation event marked for review.

+4 more

NEW DETECTION Abusive Behavior Detection for AI Models

NEW INTEGRATION Grok Bot Inference Hooks

NEW INTEGRATION Agent Guard on Amazon Bedrock

NEW INTEGRATION Alibaba Cloud AI Asset Discovery and Policy Enforcement

NEW INTEGRATION Red Teaming Agent Models — OpenAI, Anthropic, Azure OpenAI

WAS THIS USEFUL? 👍 👎

Fastly

Feedback

CLI v17.1.0 (2026-11-12)

BLOG ↗

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

Fastly CLI adds IAM authentication support for service backends

WAS THIS USEFUL?

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.9 NOTES ↗ RANK

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

> Infisical's v0.167.9 release adds a Jenkins secret sync integration, JWKS public key publishing from the KMS, configurable CRL controls for internal CAs, and a substantial expansion of agent-vault's MCP and service capabilities alongside smaller PAM and alerting improvements.

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

01 **Agent-vault MCP and service expansion** NEW 80

Agent-vault gains new capabilities: it can limit an inline MCP server to a subset of its tools, adds `PostgreSQL`, `MySQL`, and `SSH` services, adds hooks on services, and adds service discovery, service descriptions, and a built-in skill. It also adds MCP templates for six more providers, makes agent-vault services project resources that can be added to access bundles, and shows request counts in session logs.

— Names many concrete additions but no exact commands or config keys v0.167.9

02 **viewSecretValue parameter on get-secret-by-ID API** NEW 75

Adds a `viewSecretValue` parameter to the get-secret-by-ID API, allowing the decrypted value to be included in the response, e.g. `GET /api/v1/secrets/<secret-id>?viewSecretValue=true`.

Retrieve a secret by ID and include its decrypted value in the response using the new parameter.

```
$ GET /api/v1/secrets/<secret-id>?viewSecretValue=true
```

— Named API parameter with a runnable example request v0.167.9

THINNER COVERAGE BELOW

03 **CRL validity and on-demand regeneration for internal CAs** NEW 50

Configure CRL validity per internal CA and regenerate CRLs on demand.

— Clear capability but no config key or command named v0.167.9

04 **JWKS publishing for KMS sign-verify keys** NEW 45

Publishes sign-verify keys' public keys as a JWKS endpoint from the KMS.

— Describes the mechanism but gives no endpoint path v0.167.9

05 **Jenkins secret sync integration** NEW 40

Adds a Jenkins app connection and secret sync integration via feat(secret-sync).

— Names the integration but no setup steps given v0.167.9

06 **Configurable bundle method for Cloudflare certificate sync** IMPROVED 40

Adds a configurable bundle method to the Cloudflare certificate sync.

— Names the feature but no options or config detail v0.167.9

07 **'My Accounts' folder in PAM accounts page** NEW 35

Adds a virtual My Accounts folder to the PAM accounts page.

— Clear UI location but minimal explanation of behaviour v0.167.9

08 **Secret rotation failure alerts** NEW 30

Sends secret rotation failure alerts through the shared alerting module.

— States what it does but no configuration detail v0.167.9

09 **Blind index for secret scanning findings** IMPROVED 25

Implements a blind index for secret scanning findings.

— Bare mention with no mechanism or scope v0.167.9

WAS THIS USEFUL?  

Truffle Security trufflehog

Feedback ⓘ v3.99.9

NOTES ↗ CODE ↗ RANK

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

> TruffleHog v3.99.9 speeds up GitHub comment scanning with a new GraphQL mode, adds a Serval API key detector, introduces per-detector execution timing, and mentions a new generic source type.

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

01 **GraphQL-based GitHub comment scanning** NEW 80

The `github` subcommand gains a `--use-graphql` flag that scans issue and PR comments via the GitHub GraphQL API instead of REST, speeding up scans on large repositories. It requires a token and must be combined with `--issue-comments` or `--pr-comments`.

Speed up comment scanning on a large GitHub org by switching to the GraphQL API instead of the REST API.

```
$ trufflehog github --org=myorg --token=$GITHUB_TOKEN --pr-comments --use-graphql --results=verified
```

– Exact flag and required companion flags with runnable command v3.99.9

THINNER COVERAGE BELOW

02 **Per-detector execution timing** IMPROVED 50

TruffleHog now times `detector_execution_duration` per individual detector run, enabling finer-grained performance visibility across detectors.

– Names the metric but no access/consumption detail v3.99.9

03 **Serval API key detector** NEW 40

TruffleHog adds a detector to classify and validate Serval API keys found during scans.

– Names the detector but no usage detail v3.99.9

04 **Generic source type added** NEW 15

TruffleHog adds a `SOURCE_TYPE_GENERIC` source type, per the release summary, though no further detail on its usage is given.

– Only named in summary, no explanation of behavior v3.99.9

WAS THIS USEFUL?  

GitGuardian SaaS

Feedback ⓘ

Catch secret theft on your developers' machines with a Kubeconfig honeytoken (2026-11-12)

BLOG ↗ RANK

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

> GitGuardian SaaS adds a Kubeconfig honeytoken to catch secret theft directly from developer machines, alongside a new secret-revoke API endpoint, expanded Public API filtering and agents-analysis fields, login lockout protection, and federated repository scanning for JFrog.

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

01 Kubeconfig honeytoken for developer endpoints NEW

95

Plant a Kubeconfig honeytoken on developer machines with `ggshield honeytoken plant --type kubeconfig`, which writes a decoy into `~/.kube/config` and fires an alert with IP, country, user agent, and machine name on first use. Requires `ggshield` 1.56+ and the `honeytokens:write` scope.

Add a Kubeconfig honeytoken to a developer machine's fleet job alongside an existing AWS key decoy, so any credential-stealing malware triggers a named-machine alert.

```
$ ggshield honeytoken plant --type kubeconfig
```

— Exact command, file path, version and scope requirement given

Catch secret theft on your developers' …

02 Secret revocation API endpoint NEW

90

Revoke a detected secret via `POST /v1/secrets/{secret_id}/revoke`, passing provider credentials in the `credentials` body field to override workspace defaults. Requires the `secrets:write` scope.

Revoke a compromised secret found by GitGuardian, supplying provider credentials so the revocation uses them instead of workspace defaults.

```
$ curl -X POST
https://api.gitguardian.com/v1/secrets/secret_123/revoke \
  -H 'Authorization: Token <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"credentials": {"api_key": "<provider_api_key>"}}'
```

— Runnable endpoint with method, path and required scope

Catch secret theft on your developers' …

---

03 **Check-run filtering by commit, PR and source** IMPROVED 80

`GET /v1/check-runs` can now be filtered by `commit_sha`, `pull_request_number`, or `source_id` in addition to `incident_id`, and each check-run response now also includes the repository's full name and source ID.

Find all check-runs for a specific pull request even when its findings have no incident yet, to audit CI scan results in a pipeline.

```
$ curl 'https://api.gitguardian.com/v1/check-runs?
pull_request_number=42&source_id=789' \
  -H 'Authorization: Token <api_key>'
```

— Endpoint and query parameters named with a runnable example

Catch secret theft on your developers' …

---

04 **Login lockout after repeated failures** NEW 80

Password login now locks out for six minutes after ten failed attempts for the same email and IP; the login endpoint returns HTTP 429 with a Retry-After header and writes a `password_login_locked_out` entry to the audit log.

— Exact thresholds, status code and audit log key named

Catch secret theft on your developers' …

---

05 **Agents analysis fields on incidents** NEW 70

Public incident list, detail, and CSV export now carry agents analysis fields `company_related_verdict` (values: `related`, `unclear`, `unrelated`) and `agents_stage`, both filterable, so a SIEM or ticketing tool can reproduce dashboard views.

— Field names and values given but no sample call shown

Catch secret theft on your developers' …

THINNER COVERAGE BELOW

06 **Source owner filter in Public Monitoring** NEW 40

Filter public incidents by source owner — a GitHub organization or user, or a Docker Hub namespace — in Public Monitoring.

— Names the filter but no endpoint or exact UI steps

Catch secret theft on your developers' …

07 **'Synced from' filter on Teams page** NEW 35

Team Sync gains a Synced from filter on the Teams page to distinguish provider-synced teams from manually created ones.

— Names exact UI location but lacks further mechanism

Catch secret theft on your developers' …

08 **Federated repository scanning for JFrog** IMPROVED 30

JFrog Container Registry integration now scans federated Docker repositories alongside local ones.

— States the change with no further configuration detail

Catch secret theft on your developers' …

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine

Feedback ⓘ

Detection Engine Updates Version 2.178 (2026-11-11) BLOG ↗ RANK

GitGuardian Detection Engine scans code repositories and developer workflows to automatically detect exposed secrets, API keys, and credentials in real time.

> GitGuardian Detection Engine rolled out dozens of new secret detectors across Google, AWS, Supabase, and more than a dozen other providers, improved several existing detectors (Vault host reporting, Bedrock long-term keys, SSH key validation against custom hosts, Azure subscription key recall), and renamed two Google detectors in a breaking change.

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

01 **SSH private key validation against custom hosts** IMPROVED

75

Private key checkers `private_key_rsa`, `base64_private_key_rsa`, `private_key_elliptic`, `base64_private_key_elliptic`, `private_key_openssh`, `base64_private_key_openssh`, `private_key_generic`, and `base64_private_key_generic` now validate SSH keys against a custom host, such as a self-hosted GitLab or GitHub Enterprise instance.

— Names all eight detectors and the mechanism, but no config example

Detection Engine Updates Version 2.176 …

02 **Google detector renames break existing workflows** BREAKING

70

The Google Cloud Express API Key detector is now named Google Gemini API Key; the former Google Gemini API Key detector is now named Google Bard Session Cookie. Any workflow or filter referencing these detectors by their old names must be updated.

— Exact old/new names given with clear migration instruction

Detection Engine Updates Version 2.178 …

03 **New detectors for Supabase, Google OAuth2, and Food Safety Korea**

NEW

60

Adds the Supabase Personal Access Token detector covering classic, scoped, and `sbp_v0_` variants; the Google OAuth2 Refresh Token detector covering refresh tokens alongside their client ID and secret, including `gcloud` application default credentials; and the Food Safety Korea API Key detector.

— Names detectors and variants but purely passive detection

Detection Engine Updates Version 2.176 …

04 **23 new analyzers across DeepSeek, Cloudflare, GitLab and more**

NEW

60

Adds analyzers for `DeepSeek` API keys (reporting account balance, currency, availability, and usable models), Groq API keys (reporting usable models), Telegram bot tokens (reporting bot username and name), `GitLab` self-managed personal access tokens, `GitLab` runner authentication tokens, New Relic admin keys, New Relic REST API keys, New Relic user keys, Amp API tokens, Sourcegraph access tokens, `LlamaCloud` API keys, Azure Entra app secrets, Gemfury deploy, push, and full-access tokens (with and without username variants), Mercado Pago access tokens, Mercado Pago OAuth credentials, Cloudflare API tokens, Cloudflare R2 tokens, Cloudflare global API keys, Facebook access tokens, and Facebook app keys.

— Enumerates every analyzer but no further mechanism per item

Detection Engine Updates Version 2.176 …

THINNER COVERAGE BELOW

05 **New Google credential detectors** NEW

55

Adds the Google OAuth2 Refresh Token Without Client detector for refresh tokens found without their client ID, plus base64-variant detectors `base64_googlecloud_access_token`, `base64_google_oauth2`, `base64_google_oauth2_refresh_token`, the `googlecloud_storage_hmac_key` detector for Google Cloud Storage HMAC keys, and `base64_googleaiza` for base64-encoded Google API keys.

— Names six specific detectors but no usage mechanism beyond detection

Detection Engine Updates Version 2.178 …

06 **Vault detectors report instance host** IMPROVED

55

HashiCorp Vault token, AppRole, and unseal key detectors now also report the host of the Vault instance when it appears next to the credentials.

— Clear before/after behavior but no config surface

Detection Engine Updates Version 2.178 …

07 **New AWS credential detectors** NEW

40

Adds dedicated detectors for AWS AppSync API keys and AWS CodeArtifact authorization tokens.

— Thin one-line descriptions of two new detectors

Detection Engine Updates Version 2.178 …

08 **Bedrock long-term key detection expanded** IMPROVED

40

AWS Bedrock long-term API key detection now covers keys created outside the console.

— States scope expansion with no further mechanism

Detection Engine Updates Version 2.178 …

---

09 **Azure subscription key detector recall improved** IMPROVED

20

Updates the `azure_subscription_key` detector to improve recall.

— Bare statement with no detail on what changed Detection Engine Updates Version 2.176 …

└──▷ **BREAKING ON UPGRADE**

! The Google Cloud Express API Key detector is now named Google Gemini API Key ; any workflow or filter referencing Google Cloud Express API Key by name must be updated.

! The former Google Gemini API Key detector is now named Google Bard Session Cookie ; any workflow or filter referencing Google Gemini API Key as a session cookie detector must be updated.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

cyclonedx-gradle-plugin

Feedback ⓘ

cyclonedx-gradle-plugin-3.6.0 NOTES ↗ RANK

Cyclonedx-gradle-plugin creates CycloneDX Software Bill of Materials (SBOM) from Gradle projects

> cyclonedx-gradle-plugin 3.6.0 adds a build-safety flag to fail SBOM generation when dependency resolution is incomplete.

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

01 **Fail build on unresolved dependencies** NEW 85

Setting `failOnUnresolvedDependencies = true` on a `cyclonedxDirectBom` task causes the build to fail if a selected configuration has an unresolved dependency or artifact, reporting the affected dependencies and configurations. Defaults to `false`.

Catch unresolved dependencies during SBOM generation in CI so a broken build never silently produces an incomplete SBOM.

kotlin

```kotlin
tasks.named<CyclonedxDirectTask>("cyclonedxDirectBom") {
    failOnUnresolvedDependencies = true
}
```

— Exact config key, default, and behaviour all given. cyclonedx-gradle-plugin-3.6.0

WAS THIS USEFUL? 👍 👎

Endor Labs Feedback ⓘ

Endor Labs + OpenAI: Strengthening the Defense Factory With Accelerated Remediation | En

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 adds an integration with the OpenAI Daybreak Defense Factory, routing reachability analysis, SAST, SCA, and container findings into Codex for automated fix implementation.

NEW INTEGRATION Endor Labs now connects to Codex Security via the OpenAI Daybreak Defense Factory, sending each finding with full-stack reachability, risk, impact, and remediation guidance so Codex can implement and test fixes.

An AI SAST agent (AURI) gives Codex evidence and remediation guidance for vulnerabilities in first-party source code.

An SCA Remediation agent gives Codex reachability evidence and remediation guidance for open source vulnerabilities across direct and transitive dependencies.

Endor Labs now provides Codex with reachability evidence and guidance to remediate vulnerabilities in container base images.

Customer-defined policies let teams prioritize which Endor Labs findings the Daybreak Defense Factory acts on first, based on organization-specific risk criteria.

WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ The Token Cost of Shell Commands (2026-11-12)

BLOG ↗ 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 introduced Boost, a free CLI wrapper in public preview that compresses noisy shell command output before it reaches AI coding agents, cutting token costs in agentic workflows.

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

01 **Boost CLI wrapper for AI agent token compression** NEW

80

JFrog Boost is a free CLI wrapper available at `boost.jfrog.com` that compresses noisy shell command output from tools like `ls`, `find`, and `grep` before it reaches AI coding agents such as Cursor, Claude Code, or GitHub Copilot. Truncated output is paired with an explicit `[System]` note and a `cat_file` recovery hatch so agents know results were partial and can request the full output instead of hallucinating.

Wrap grep in a Python agent tool to enforce output limits and tell the agent when results were truncated, preventing silent context poisoning.

```python
import subprocess

def safe_grep(search_term: str, directory: str = "."):
    cmd = [
        "grep", "-r", "-n", "-I",
        "--exclude-dir=.git",
        "--exclude-dir=node_modules",
        "--exclude-dir=dist",
        "-e", search_term,
        "--", directory,
    ]
    result = subprocess.run(cmd, capture_output=True, text=True,
check=False)
    output = result.stdout
    char_limit = 8000
    if len(output) > char_limit:
        return (
            f"{output[:char_limit]}\n\n"
            "[System: Output truncated. Use `cat_file` for
specific files.]"
        )
    return output
```

— Names install URL, affected commands, recovery mechanism, and shows working code

The Token Cost of Shell Commands (2026-…

WAS THIS USEFUL?  

Cloudsmith

Feedback ⓘ

Host, version and govern AI skills and Claude Code plugins (2026-11-12) BLOG ↗

RANK

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

> Cloudsmith now hosts, versions, and governs AI skills and Claude Code plugins using the same repository workflows as its other package types.

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

01 Hosting for AI skills and Claude Code plugins NEW

53

Cloudsmith repositories can publish and version AI skills and Claude Code plugins, reusing the existing artifact management workflows already in place for other package types. The web app adds a repository view for AI Skills repositories, a package Overview and Files tab for individual agent skills, and a Configure upstream form for setting up an AI Skills upstream.

.
.
.
.

— Describes UI surfaces but no API/CLI commands or config keys

Host, version and govern AI skills and …

WAS THIS USEFUL? 👍 👎

cyclonedx-core-java

Feedback ⓘ cyclonedx-core-java-13.5.0

NOTES ↗ RANK

Cyclonedx-core-java is a cyclonedx SBOM Model and Utils for Creating and Validating BOMs

> cyclonedx-core-java 13.5.0 adds the ability to serialize a BOM directly to binary output formats without going through a string intermediate.

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

01 **Direct byte[] and OutputStream serialization** NEW 55

BOMs can now be serialized directly to `byte[]` or `OutputStream`, removing the need to convert from a string intermediate.

— Names exact output types but no code example given cyclonedx-core-java-13.5.0

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

feroxbuster

Feedback ⓘ v2.14.0 NOTES ↗ RANK

Feroxbuster is a fast, simple, recursive content discovery tool written in Rust.

> Feroxbuster v2.14.0 adds support for combining multiple wordlists in a single scan, deduplicating entries across them automatically.

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

01 **Multiple wordlists in one scan** NEW 80

Pass multiple wordlists to a single scan by supplying `-w` more than once; entries are deduplicated across all source files, e.g. `feroxbuster -u http://target.example.com -w /usr/share/wordlists/dirb/common.txt -w /usr/share/wordlists/api-paths.txt`.

Combine a general directory list with a targeted API-path list in one scan without preprocessing the files manually.

```
$ feroxbuster -u http://target.example.com -w /usr/share/wordlists/dirb/common.txt -w /usr/share/wordlists/api-paths.txt
```

— Names exact flag and behavior with runnable command example. v2.14.0

WAS THIS USEFUL? 👍 👎

Acunetix 360

Feedback ⓘ v26.11.0 – 12 November 2026 (2026-11-12)

BLOG ↗ RANK

Acunetix 360 scans web applications and APIs to identify security vulnerabilities through automated testing.

> Acunetix 360's v26.11.0 release adds a way to flag internet-exposed targets and enriches the context attached to vulnerability findings.

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

01 **Internet-exposed target setting** NEW 35

Acunetix 360 adds a setting on targets to flag which ones are reachable from the public internet.

– Names the setting but not UI path or mechanism v26.11.0 – 12 November 2026 (2026-11-12)

02 **Enriched vulnerability data** IMPROVED 9

Acunetix 360 now enriches vulnerability data with additional context.

– No detail on what context was added v26.11.0 – 12 November 2026 (2026-11-12)

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite

Feedback ⓘ

DAST 2026.11.1 (2026-11-12) BLOG ↗ RANK

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

> Burp Suite DAST introduced a beta GraphQL API endpoint for querying scan data.

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

01 **Beta GraphQL API for scan data** NEW 40

Burp Suite DAST adds a new beta GraphQL API endpoint that allows users to access and query scan data programmatically.

— No endpoint path, query schema, or auth details given. DAST 2026.11.1 (2026-11-12)

WAS THIS USEFUL? 👍 👎

◆ Code Security

GitHub

Feedback ⓘ

Public monitoring alerts now have direct links (2026-11-12) BLOG ↗ RANK

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

> GitHub added direct, shareable URLs for individual public monitoring alerts in the enterprise view.

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

01 **Direct links for public monitoring alerts** NEW 50

In the enterprise Public monitoring view, clicking an alert number now opens a page with its own direct URL, which can be bookmarked or shared with teammates who have access, instead of requiring navigation through the full list each time.

Share a specific public monitoring alert with a teammate without navigating them through the full enterprise view.

> 📍 In your enterprise, go to the Public monitoring view, click an alert number to open it, then copy the URL from the browser address bar to share with any teammate who has access.

— Clear UI path given, but no API or flag detail.

Public monitoring alerts now have direc…

WAS THIS USEFUL? 👍 👎

GitLab

Feedback ⓘ

Turn Jira tickets into GitLab merge requests with Duo Agent Platform (2026-11-12)

BLOG ↗ RANK

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

> GitLab's Duo Agent Platform can now turn a Jira ticket directly into a draft merge request, via a new custom Jira agent that uses the Atlassian MCP server and hands work off to Developer Flow.

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

01 **Custom Jira agent with Atlassian MCP server** NEW 95

A custom Jira agent can be built at AI > Agents > New agent, attached to the Atlassian MCP server for tooling, and enabled per-project through the AI Catalog. The agent can search, summarize and comment on Jira tickets using the user's own Atlassian login, including calling tools like `searchJiraIssuesUsingJql` with user-written JQL and approving tool calls inside chat; the Atlassian MCP server is supported in the AI Catalog as an experiment.

Find your high-priority open Jira tickets without leaving GitLab, using the custom Jira agent and JQL.

> 📍 Steps: 1. Open your project in GitLab and select Add new chat in the Duo sidebar. 2. Choose Jira from the agent dropdown and confirm the Agentic toggle is on. 3. Type: Which of my open Jira issues in KAN are High or Highest priority? 4. Approve the searchJiraIssuesUsingJql tool call when prompted; the agent runs the query and returns a summary table.

— Names exact menu paths, tool and JQL usage with a runnable example.

Turn Jira tickets into GitLab merge req…

02 **Developer Flow creates merge requests from Jira tickets** NEW 90

Using `/goal` or the + > Goal button in Duo Agentic Chat, a user can hand a Jira ticket to Developer Flow, which runs as a CI pipeline, creates a branch, writes the change and tests, and opens a draft merge request without further approval prompts.

Trigger Developer Flow to implement a Jira ticket, create a branch, and open a draft MR — all from Agentic Chat.

📍 Steps: 1. In your GitLab project, open Duo Agentic Chat and select `Jira` from the agent dropdown. 2. Make sure the Agentic toggle is on. 3. Select + › Goal (or start your message with `/goal`) and type: Implement KAN-123: add a `/healthz` endpoint returning 200 and {"status":"ok"}, with a test. Create a branch named KAN-123-healthz and open a DRAFT merge request. The MR description must start with: Resolves KAN-123 - `https://your-site.atlassian.net/browse/KAN-123` 4. Monitor progress at AI › Sessions; the MR appears when the session shows Finished.

— Exact command and pipeline mechanism given, runnable end to end.

Turn Jira tickets into GitLab merge req…

03 **Jira issue integration for merge requests** NEW 80

Connecting Jira to a GitLab project via `Settings > Integrations > Jira issues` makes any merge request that mentions a ticket key (e.g. `KAN-123`) post a comment back to Jira and surface the MR on the ticket; setup requires a Web URL, Basic authentication with Atlassian email and API token, and enabling the Commit/Merge request trigger with comments.

Enable the Jira integration on a project so MRs that mention a ticket key automatically post a comment back to Jira.

📍 Steps: 1. In GitLab, open the project and go to Settings › Integrations › Jira issues. 2. Select Active. 3. Set Web URL to your Jira site, e.g. `https://your-site.atlassian.net`. 4. Set Authentication to Basic and enter your Atlassian email and API token. 5. Under Trigger, select Commit and Merge request, and tick Enable comments. 6. Click Test settings, then Save changes.

— Exact settings path and setup steps, but describes existing-style integration flow.

Turn Jira tickets into GitLab merge req…

THINNER COVERAGE BELOW

04 **Per-user Atlassian account linking for agent comments** NEW 55

Each user can connect their own Atlassian account once under `AI > MCP servers` so that Jira comments posted by the custom agent appear under that person's name rather than under a bot account.

— Names the config location but gives no step-by-step setup.

Turn Jira tickets into GitLab merge req…

WAS THIS USEFUL? 👍 👎

◆ Application Security Posture

Cycode Feedback ⓘ

wfx CI secret stealer: How to check if a stolen GitHub token probed your repositories (2

BLOG ↗

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

Cycode adds three Threat Intelligence dashboard items to detect wfx-docs CI secret-stealer activity across connected GitHub repositories, including private ones.

Three new items on the Threat Intelligence dashboard query the Context Intelligence Graph across connected GitHub repositories, including private ones, to surface probe pull requests, probe branches, and unreverted exfiltration workflow steps.

NEW DETECTION The Probe pull requests detection persists results after the pull request is closed and the branch is deleted, keeping a record even when attackers clean up.

NEW DETECTION The 'Workflow steps that still contain the script' item matches on environment variables `WF_SECRETS` and `WF_CALLBACK` and reported callback server addresses, covering branches Cycode scans.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

karton

Feedback ⓘ v6.0.0 NOTES ↗ RANK

Karton is a distributed malware processing framework based on Python, Redis and S3.

> Karton v6.0.0 introduces Karton Gateway, a new HTTP backend layer with JWT authentication, and tightens version requirements for services and Python.

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

01 **Karton Gateway HTTP backend with JWT authentication** NEW

65

Karton adds Karton Gateway, an HTTP server backend that sits between Redis and all services to enforce permissions, increase observability, and simplify deployment in corporate environments — now the primary backend going forward. Services authenticate using JWT-based tokens, with a CLI included for generating them.

– Describes purpose and mechanism but no CLI command shown v6.0.0

THINNER COVERAGE BELOW

02 **Backend creation via get_backend factory** BREAKING

55

Backend instances must now be created using the `get_backend` factory function instead of being constructed directly.

– Names exact function but no usage example v6.0.0

03 **Dropped support for Karton 4.x and older Python** BREAKING

55

Karton 4.x and older are no longer supported — services must be on karton-core 5.0.0 before upgrading — and Python 3.12 or newer is now required, as older Python versions will not work.

– Clear version requirements but no migration steps given v6.0.0

└──▷ **BREAKING ON UPGRADE**

! Karton 4.x and older are no longer supported — services must be on karton-core 5.0.0 before upgrading.

! Python 3.12 or newer is now required; older Python versions will not work.

! Backend instances must now be created via the `get_backend` factory function instead of being constructed directly.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

hrtng

Feedback ⓘ v3.9.118 NOTES ↗ RANK

> hrtng added IDA 9.5 support and case-insensitive API name hash scanning, while disabling the AVX lifter for IDA 9.5 and later.

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

01 **Case-insensitive API hash scanning** IMPROVED 45

Scan for API name hashes now supports case-insensitive modes (lower and upper) to match obfuscated binaries that hash API names in non-standard casing.

– Explains purpose and modes but no exact UI/flag v3.9.118

02 **IDA 9.5 support** NEW 30

hrtng now supports IDA 9.5.

– Bare statement of version support, no mechanism given v3.9.118

03 **AVX lifter disabled on IDA 9.5+** BREAKING 30

The AVX lifter is disabled for IDA 9.5 and later.

– States the change but no workaround or reasoning v3.9.118

└──▷ **BREAKING ON UPGRADE**

! The AVX lifter is disabled for IDA 9.5 and later.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

cilium-cli

Feedback ⓘ v0.20.3 NOTES ↗ RANK

CLI to install, manage & troubleshoot Kubernetes clusters running Cilium

> cilium-cli v0.20.3 adds a flag to skip the host-netns DaemonSet in connectivity tests and improves sysdump diagnostics collection.

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

01 Sysdump collection fixes and Gateway API discovery IMPROVED

60

`cilium sysdump` now collects CiliumCIDRGroups again, restoring collection that had been broken since v0.18.4 and previously left an empty `ciliumcidrgroups-<ts>.yaml` in the bundle. It also now discovers Gateway API versions when collecting diagnostic data.

— Names exact fixed file and regression version v0.20.3

THINNER COVERAGE BELOW

02 Skip host-netns DaemonSet in connectivity tests NEW

50

A new flag on `cilium connectivity test` allows skipping creation of the host-netns DaemonSet during test runs.

— Names the command but not the flag itself v0.20.3

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.49.0 NOTES ↗ CODE ↗ RANK

This repo hosts Prowler OSS (CLI and Local Server). Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment for Humans and AI Agents.

> Prowler 5.49.0 introduces emerging-threat signature matching for AWS, Azure, and Google Cloud compute instances, adds two new AWS checks, and extends the Prowler Cloud API with threat-based filtering.

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

01 **Emerging threat signature detection for compute instances**

NEW 88

Supply threat signatures per scan with the `threat_signatures` configuration key for AWS, Azure, and Google Cloud, matching compute instances against affected image publishers, name patterns, per-branch version ranges, exposed ports, impact, and remediation. Three new checks — `ec2_instance_emerging_threat_image` (AWS), `vm_emerging_threat_image` (Azure), and `compute_instance_emerging_threat_image` (Google Cloud) — flag instances running images matched by a loaded threat signature, storing matched threats in the new `Check_Report.raw` payload attached to each finding. On AWS, the engine clears a match when Amazon Inspector finds no active CVE finding for the threat, using the new `inspector2:SearchVulnerabilities` permission added to the Prowler additions policy and the CloudFormation scan role.

— Names config key, three checks, payload field and IAM permission 5.49.0

02 Threat-based issue filtering in Prowler Cloud API NEW

75

Filter issues by threat in Prowler Cloud with `filter[threat]` or search threat names and CVE IDs with `filter[search]` on `GET /api/v1/issues` and `GET /api/v1/issues/overview`.

List all Prowler Cloud issues tied to a specific emerging threat to see every affected host in one call.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://api.prowler.com/api/v1/issues?filter[threat]=my-threat-
001'
```

— Names exact endpoints and filter params with a runnable example 5.49.0

THINNER COVERAGE BELOW

03 OAM sink and Organizations delegated-admin checks NEW

55

New check `oam_sink_not_publicly_accessible` reports CloudWatch Observability Access Manager sinks whose resource policy allows any AWS account to link to them. New check `organizations_security_services_delegated_admin_not_management_account` verifies that security services with trusted access in the organization use a delegated administrator rather than the management account.

— Names both checks but no usage detail beyond description 5.49.0

WAS THIS USEFUL? 👍 👎

◆ SaaS Security Posture

Nudge Security

Feedback ⓘ Detect and resolve breached passwords

NOTES ↗ RANK

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

> Nudge Security introduced breached password detection and resolution, using the browser extension to catch compromised credentials at login and surfacing findings across dashboards, accounts, and user profiles.

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

01 **Breached password detection and resolution** NEW 76

The browser extension checks employee credentials against known data-breach datasets in real time at login, and can be configured via the Browser Extension Settings to enable breached password detection and an in-browser nudge that prompts employees to change exposed passwords. Findings are generated automatically when a breach is detected and surfaced via a new password risk findings card on the Identity & location dashboard, as well as on individual account and user pages.

Turn on breached password detection so employees are nudged in-browser the moment they log in with a compromised credential.

> 📍 In the Nudge Security console, go to Browser Extension Settings and enable breached password detection and the browser nudge setting for breached passwords.

– Names exact settings and dashboard but no limits or data sources

Detect and resolve breached passwords

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Microsoft Sentinel, upgraded: NOTES ↗ RANK

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

> Plerion upgraded its Microsoft Sentinel integration with richer incident handling and added a proactive warning before Sentinel client secret expiry.

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

01 **Enhanced Microsoft Sentinel incident handling** IMPROVED

53

Plerion's Microsoft Sentinel integration now lets workflows set an incident owner from a Plerion user or an email address, add labels to incidents, and close incidents with a classification that reflects why the alert closed in Plerion. Findings and vulnerabilities are now sent to Sentinel with their real severity rather than a mapped value.

— Describes behavior changes but no exact fields or UI steps

Microsoft Sentinel, upgraded:

02 **Sentinel client secret expiry warning** NEW

33

Plerion now warns you 14 days before your Microsoft Sentinel client secret expires.

— Clear behavior but no detail on where warning appears Microsoft Sentinel, upgraded:

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

AI Function task support in Declarative Automation Bundles (Beta) (2026-11-13)

BLOG ↗ RANK

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

> Databricks expanded Genie One substantially this window — adding AI-built Artifacts/slide decks, a generally available macOS desktop app, a Chrome new-tab extension, mobile deep links, an agent connectors section, and cross-user scheduled task emails — while also shipping ABAC DENY policies for AI securables, generally available OAuth on-behalf-of for service principals, AI Function tasks in Declarative Automation Bundles, GitHub Copilot configuration via Unity Gateway, and several AI/BI dashboard improvements.

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

01 GitHub Copilot configuration via Unity Gateway NEW

90

Unity Gateway now configures GitHub Copilot in Visual Studio Code alongside GitHub Copilot CLI — including models, default model, HTTP headers, and tracing — using `ug configure --agents copilot`, and lets you enable or disable Copilot in VS Code from the coding agent configuration.

Set up GitHub Copilot in VS Code and Copilot CLI together through Unity Gateway, specifying models and tracing in one pass.

```
$ ug configure --agents copilot
```

— Exact runnable command and scope of settings given.

Connect GitHub Copilot CLI and GitHub C…

02 Artifacts and slide decks in Genie One NEW

85

Genie One (Beta) can now build Artifacts — AI-generated slide decks and custom visual pages built from your data — that persist in the workspace, can be edited or updated via chat, and store a data snapshot so viewers see the data without needing access to source tables. Enable it via the Artifacts preview setting on the Previews page in workspace admin settings.

Enable the Artifacts Beta so your workspace users can ask Genie One to build and share slide decks from chat.

> 📍 1. Sign in as a workspace admin. 2. Go to Settings › Previews. 3. Find `Artifacts` and turn it on.

— Mechanism and enable steps given, but still Beta-scope.

`Artifacts and slides in Genie One (Beta…`

---

03 **ABAC DENY policies for AI securables** NEW — 80

DENY policies can now be attached directly to individual objects using the `ON` clause, covering tables, volumes, functions, models, secrets, model services, model provider services, MCP services, agent services, and skills. DENY policies can also be applied via the `FOR` clause to model services, model provider services, MCP services, agent services, and skills. This is in Beta.

— Names exact clauses and securable types but no full command shown.

`ABAC DENY policies support AI securable…`

---

04 **OAuth on-behalf-of for service principals** IMPROVED — 70

OAuth on-behalf-of (OBO) for service principals is now generally available and on by default, letting you use the standard user-to-machine (U2M) OAuth flow to act as a service principal you manage without creating an OAuth secret — this requires the Service Principal Manager role on the service principal.

— Clear mechanism and role requirement, no exact command given.

`OAuth on-behalf-of for service principa…`

---

05 **Pivot table, filter, and variable improvements in AI/BI Dashboards**

IMPROVED — 65

AI/BI Dashboards now let viewers sort pivot tables by row and column, let authors style pivot table headers separately per row, column, and value field, let pivot table values wrap text to adjust height, and let authors resize the divider between columns and values. Authors can turn off associative filtering between dashboard filters, and actions on dashboard variables now support undo and redo. Genie Code can write custom calculations that reference other custom calculations in the same dataset, and Ask Genie now appears only on widgets Genie can use as context, no longer on images, filters, or older text boxes.

— Each change named but no exact UI path or command.

`Dashboard enhancements and bug fixes (2…`

THINNER COVERAGE BELOW

06 **Connectors section in Genie Agent settings** NEW 55

Genie Agent settings now include a dedicated Connectors section (Beta) for configuring connectors so a Genie Agent can query external data sources alongside Databricks data.

Add a connector to a Genie Agent when you want it to query external data sources alongside your Databricks data.

> 📍 In the Databricks UI, open your Genie Agent, go to Configure settings, locate the new Connectors section, and select Add connector.

— UI navigation path given, mechanism lightly described.

Connectors section in Genie Agent setti…

07 **Scheduled task emails to other users in Genie One** IMPROVED 55

Scheduled tasks in Genie One can now send result emails to other users in your organization, not only the task owner; add recipients under `Send email to` when configuring a scheduled task.

Share scheduled Genie One task results with teammates automatically, without requiring them to set up their own tasks.

> 📍 In Genie One, open a scheduled task, then under `Send email to`, add other users in your organization as recipients.

— Exact UI steps given for a narrow addition. Scheduled tasks can email other users i…

08 **Genie links open in mobile app** IMPROVED 35

Links to Genie One chats and published dashboards now open directly in the Genie mobile app on iOS and Android instead of in a browser.

— Clear what changed but thin, no further mechanism.

Links open in the Genie mobile app (202…

09 **AI Function tasks in Automation Bundles** NEW 35

Declarative Automation Bundles now support AI Function tasks (Beta), enabling AI-driven steps to be defined and orchestrated as bundle resources.

— Names the capability but no config keys or example shown.

AI Function task support in Declarative…

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ AI Observability App (Apps) (2026-11-13)

BLOG ↗ RANK

Sumo Logic is a cloud-native platform that provides real-time log management, security analytics, and observability to help organizations monitor and secure their applications and infrastructure.

> Sumo Logic introduced an AI Observability app for monitoring LLM applications through OpenTelemetry trace data.

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

01 **AI Observability app for LLM applications** NEW 75

A new AI Observability app (Limited Availability) ships with preconfigured dashboards and monitors to track health, performance, cost, and behavior of AI and LLM applications, using OpenTelemetry trace spans enriched with `gen_ai.*` GenAI Semantic Convention attributes. Instrumentation guides are provided for OpenLLMetry and OpenTelemetry.

Route LLM trace spans with GenAI Semantic Convention attributes to Sumo Logic so the AI Observability app can populate its dashboards and monitors.

📍 In the Sumo Logic App Catalog, search for AI Observability and install the app. Then instrument your LLM application using OpenLLMetry or the OpenTelemetry GenAI instrumentation so that trace spans include `gen_ai.*` attributes, and configure your OpenTelemetry exporter to send spans to your Sumo Logic HTTP Traces endpoint.

— Names app, attributes and setup steps but app is Limited Availability

AI Observability App (Apps) (2026-11-13)

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

rsigma

Feedback ⓘ v0.27.0 NOTES ↗

Rsigma is a complete Sigma detection engineering toolkit: parser, linter, formatter, evaluator, correlation engine, conversion framework, streaming daemon, MCP and LSP servers

RSigma adds `rsigma rule fmt` for canonical Sigma rule formatting, with CI flags, MCP and LSP integration, and a new `rsigma-parser` formatting API.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/timescale/rsigma:0.27.0
```

git

```
$ git clone --branch v0.27.0 https://github.com/timescale/rsigma.git
# already have the repo? check out this version:
$ git checkout v0.27.0
```

published under the project's GitHub account · checked 13 Nov Wrong command?

Adds `rsigma rule fmt` to rewrite detection rules, correlations, filters, and action-chain fragments into a canonical layout, with `--check`, `--diff`, `--write`, and `--stdout` flags for CI, review, and editor workflows.

The MCP server now exposes rule formatting as the `format_rules` tool, accepting inline `yaml` or a confined `path`, with `write: true` to persist results when the server runs with `--rules-dir`.

The LSP now advertises document and range formatting, delegating both to the parser so an editor, the CLI, and the library produce identical output.

Adds `format_yaml`, `format_yaml_range`, `FormatOutcome`, `RangeEdit`, and `FormatError` to `rsigma-parser` behind the default `fmt` feature; disable default features for `wasm32-unknown-unknown` builds.

NEW SUPPORTED PLATFORM ECS Linux canonicalizing pipeline

NEW SUPPORTED PLATFORM Splunk CIM canonicalizing pipeline

NEW SUPPORTED PLATFORM OCSF canonicalizing pipeline

NEW SUPPORTED PLATFORM CrowdStrike Falcon Windows canonicalizing pipeline

└──▷ BREAKING ON UPGRADE

- ! `mcp serve --http` on a non-loopback address refuses to start without `--auth-token` or `--allow-unauthenticated`.
- ! Protected `mcp.*` and `daemon.*` keys in `./rsigma.yaml` or `.rsigmarc` project config files are now ignored, and jq filters lose the `env` builtin.
- ! API silences now require an `ends_at` field and a scoped `=` matcher; `POST /api/v1/silences` returns `409` when no alert pipeline is configured.
- ! In `rsigma-mcp`, `http_router_with` and `serve_http_with` now take `HttpOptions` instead of their previous signatures.
- ! `AlertPipeline::new` now takes `SilenceLimits` as an argument.
- ! The `rsigma-parser` emitter now quotes and orders its output differently, which changes the text produced by existing formatting calls.

WAS THIS USEFUL? 👍 👎

Tenzir Feedback ⓘ

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

BLOG ↗ RANK

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

> Tenzir's Tree-sitter TQL parser gained support for several new syntax forms, improving editor tooling for the Tenzir Query Language.

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

01 **Expanded TQL syntax support in Tree-sitter grammar** IMPROVED

45

The Tree-sitter TQL parser now supports optional access on top-level fields using the `?` suffix syntax, such as `sort key?`, and also handles scientific durations, raw strings, pipeline values, and parenthesized lambdas.

— Names exact syntax forms but no broader mechanism or limits

Tree-sitter TQL v1.2.1 (2026-11-12)

NEW INTEGRATION Redis source and sink operators

UPDATED SUPPORTED PLATFORM Native source semantics for Splunk, ClickHouse, MySQL, SQL Server, KQL, and SentinelOne

UPDATED COMPLIANCE FRAMEWORK Latest upstream OCSF schemas bundled

> Tenzir Node v6.22.0: OCSF compaction and native source semantics (2026-11-12)
>
> BLOG ↗ STABLE
>
> Tenzir Node v6.22.0 adds lossless OCSF compaction via `ocsf_compact`, auto-computed OCSF uid fields, a `version` function, time zone support in `parse_time`, and automatic syslog framing detection.
>
> Compact OCSF events without data loss using the new `ocsf_compact` operator.
>
> Automatically compute `category_uid` and `type_uid` fields in `ocsf_derive` without manual lookup.
>
> Add time zone support to `parse_time` for accurate timestamp parsing across locales.
>
> Retrieve the running node version in pipelines using the new `version` function.
>
> Automatically detect syslog framing without manual configuration.
>
> +1 more

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ

Condition actions now support multiple paths (2026-11-12) BLOG ↗ RANK

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

> Tines updated Condition actions to support multiple ordered paths, letting a single action replace a chain of separate conditions.

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

01 **Multi-path support for Condition actions** IMPROVED 80

Condition actions now support up to 10 independent paths, each with its own rules and links, functioning like a switch statement so one action can replace a series of separate Condition actions. In a story you open a Condition action, use `Add path` up to 10 times, define rules per path (e.g. severity == `critical`, `high`, `low`), and connect each path to its downstream action; the event follows the first matching path in order.



Route an alert to different response workflows based on severity without chaining multiple Condition actions.

📍 In your story, open a Condition action › select `Add path` up to 10 times › define rules for each path (e.g. severity == `critical`, severity == `high`, severity == `low`) › connect each path to its downstream action. The event follows the first matching path in order.

— Names exact limit and UI steps, lacks broader rollout detail

Condition actions now support multiple …

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

XDR Data Explorer is now Explorer (2026-11-12) BLOG ↗ RANK

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

> Trend Vision One reorganized its navigation by renaming XDR Data Explorer to Explorer, and added recipient-prompted password scanning for protected attachments in Exchange Online (Inline Mode).

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

01 Password-protected attachment scanning via recipient prompt

NEW

70

In Exchange Online (Inline Mode), Trend Vision One can now request attachment passwords directly from recipients, enabling scanning of password-protected attachments even when the password is absent from the email body. Configure this in Email and Collaboration Security under Cloud Email and Collaboration Protection → Exchange Online (Inline Mode) policy, setting the action for password-protected attachments to Request password from recipient.

Configure Exchange Online (Inline Mode) policy to prompt recipients for attachment passwords instead of recording those emails as unscannable.

📍 In Email and Collaboration Security, go to Cloud Email and Collaboration Protection → Exchange Online (Inline Mode) policy, then configure the action for password-protected attachments to Request password from recipient.

— Names exact policy path and action but lacks limits/mechanism detail

Password-protected attachment support w…

THINNER COVERAGE BELOW

02 XDR Data Explorer renamed and relocated to Explorer IMPROVED

40

XDR Data Explorer has been renamed to Explorer and moved in the left navigation panel, now appearing between Dashboards and Reports and Agentic SIEM & XDR.

Find the renamed Explorer module after upgrading, since it has moved to a new position in the navigation.

📍 In the left navigation panel, go to Explorer (located between Dashboards and

> Reports and Agentic SIEM & XDR).

— Simple rename/relocation with exact navigation path given

XDR Data Explorer is now Explorer (2026…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface Management

CybelAngel

Feedback ⓘ [NEW] Custom CISO Dashboard NOTES ↗

RANK

CybelAngel is the all-in-one solution against external threats. We monitor the open, deep, and dark web to protect your assets, brand, and credentials 24/7

> CybelAngel shipped three new capabilities this window: a Custom CISO Dashboard for executive reporting, a prioritized Vulnerability Intelligence view within its Unified Threat Intelligence, and Vendor Exposure Monitoring that flags supplier-targeted attacker claims.

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

01 **Custom CISO Dashboard for executive reporting** NEW

40

A Custom CISO Dashboard lets teams build executive-ready dashboards tailored to the organization's indicators and share them with brand or legal teams, without exposing dark web or credential data.

– Describes purpose and data scope but no config or UI path [NEW] Custom CISO Dashboard

02 **Vulnerability Intelligence view in Unified TI** NEW

40

The Unified Threat Intelligence view now includes a prioritized Vulnerability Intelligence list showing only CVEs that attackers are actively exploiting, letting teams focus patching on what matters for their environment.

– Explains filtering logic but no concrete UI or API detail

[NEW] Unified TI View: Vulnerability In…

03 **Vendor Exposure Monitoring for supplier attacks** NEW

40

Vendor Exposure Monitoring connects attacker-channel signals to a customer's supplier list, letting teams assess impact and contact a supplier before an incident reaches their own organization.

– Clear mechanism described but no concrete setup steps given

[NEW] Vendor Exposure Monitoring

WAS THIS USEFUL? 👍 👎

bbot

Feedback ⓘ v3.1.0 NOTES ↗

Bbot is the recursive internet scanner for hackers.

bbot replaces massdns with blastdns 2.0, adds `--custom-nameservers`, `dns.brute_max_resolvers`, and `dns.max_concurrency`, and hardens custom module loading.

└──▷ **GET THIS VERSION** ▾

pipx

```
$ pipx install bbot==3.1.0
# already installed? upgrade to this version:
$ pipx install --force bbot==3.1.0
```

git

```
$ git clone --branch v3.1.0
https://github.com/blacklanternsecurity/bbot.git
# already have the repo? check out this version:
$ git checkout v3.1.0
```

named in the project's README · checked 13 Nov Wrong command?

Set custom upstream resolvers for a scan with `--custom-nameservers` (or `dns.nameservers` in config); an invalid value now aborts during prep instead of silently returning zero results.

Control how many brute-force resolvers a scan draws with `dns.brute_max_resolvers` (default 600, `0` for the full list); dead resolvers are dropped at startup and pacing learned in one run carries into the next.

Tune DNS concurrency with the new `dns.max_concurrency` setting (default 256) alongside a raised `dns.threads` default of 50; on large subdomain-enum scans this halved wall time versus 3.0.2.

Allow CDN/WAF port scanning (previously filtered by the removed `portfilter` module) by passing `--allow-cdn-ports`.

bbot now replaces massdns with blastdns 2.0 for DNS brute-forcing, removing the massdns dependency entirely; the brute-force path gains its own retry budget, per-query timeout, and persistent sockets.

+5 more

NEW DETECTION favicon technology detection

NEW DETECTION obfuscated JavaScript detection and unpacking

NEW DETECTION WAF bypass detection via body padding

NEW DETECTION baddns dangling delegation detection for DKIM/DMARC/ACME selectors

NEW INTEGRATION otilabs subdomain enumeration source

NEW INTEGRATION ScanMalware public scan archive subdomain source

UPDATED INTEGRATION dnsdumpster converted to authenticated API

└──▷ **BREAKING ON UPGRADE**

- ! Python 3.10 is no longer supported; `requires-python` is now `>=3.11,<3.16`.
- ! The `portfilter` module is removed; CDN/WAF port filtering now happens in core by default. To restore the old open-port behavior, pass `--allow-cdn-ports`.
- ! `config.module_dirs` is deprecated and ignored; custom module directories must be specified as a top-level `module_dirs` key in a preset. `-S` no longer implies `-y` when loading custom modules.
- ! `modules.lightfuzz.avoid_wafs` now defaults to `try_bypasses` instead of skipping WAF-tagged hosts. For the 3.0 behavior set `modules.lightfuzz.avoid_wafs: always` and pass `-em nowafpls`. The top-level `avoid_wafs` key is no longer read by lightfuzz.
- ! `dnsdumpster` now requires an API key.
- ! `waf_bypass` now targets scan targets only by default, not all discovered hosts.
- ! CPE-derived technology names are now derived from the product field of the CPE; existing output consumers keying on the old name format will see different values.
- ! In-scope events are no longer tagged `affiliate`; presets or output consumers that filter on the `affiliate` tag will behave differently.

WAS THIS USEFUL? 👍 👎

Axonius

Feedback ⓘ Axonius Release Notes 10.0.2 (2026-11-08) BLOG ↗

Axonius is a cybersecurity asset management platform that discovers, inventories, and manages all devices, users, and cloud assets across an organization's IT environment.

Axonius adds dashboard data-scope sharing, in-place table cell editing, AI memory management from documents, and duplicate-device detection in coverage analysis.

Add organizational AI memory from a document by uploading one PDF, CSV, or TXT file (up to 10 MB); Axonius extracts up to 25 memories and discards the file.

Delete up to 200 AI organizational memories at once from the Memories table.

AI Agent coverage-gap analysis now splits missing-agent devices into 'likely the same device', `weak match`, and `no match` buckets.

Coverage Check briefings, notifications, and emails now show how many unmanaged devices likely match a managed device.

Security Findings table and asset profile now include a sortable `Last Updated` date and time column.

+14 more

NEW INTEGRATION 6 new adapters

UPDATED INTEGRATION 59 updated adapters

└──▷ **BREAKING ON UPGRADE**

- ! The `Don't split source vulnerabilities into CVEs` system setting is removed; use the `Security Findings: CVE and vendor IDs` setting on each adapter instead.
- ! Creating a ticket from a Recommended Action now requires the `Security Findings: View security findings` permission.
- ! Every change to a Risk Score now requires both the `Manage Risk Scores` permission and the standard Action Center permission for the operation.

WAS THIS USEFUL? 👍 👎

◆ Vulnerability & Exposure Management

oxo

Feedback ⓘ v2.22.15 NOTES ↗ RANK

> oxo enriched its HTTP fingerprinting output with status and redirect metadata and introduced a new message type for detecting lookalike domains during scans.

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

01 **Richer HTTP fingerprint messages** IMPROVED 55

The `http` and `http.dom` fingerprint messages now include status codes, redirect chains, and page metadata, giving scans more context about discovered web endpoints.

— Names exact message types but no usage example v2.22.15

02 **Lookalike domain fingerprinting** NEW 55

oxo now emits a `v3.fingerprint.domain_name.lookalikes` message to surface lookalike domain detections during scans.

— Names exact message type, lacks mechanism detail v2.22.15

WAS THIS USEFUL? 👍 👎

◆ Patch Management

Installomator

Feedback ⓘ v10.10 NOTES ↗

Installomator is an installation script to deploy standard software on Macs

Installomator v10.10 adds 160+ new application install labels, removes 29 deprecated ones, and renames `npssantafull.sh` to `npssanta.sh`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v10.10
https://github.com/Installomator/Installomator.git
# already have the repo? check out this version:
$ git checkout v10.10
```

NEW INTEGRATION Installomator now supports installing applications via 160+ new labels including `checkpointsase`, `globalprotect62`, `globalprotect63`, `globalprotect64`, `netbird`, `protonmail`, `protonpass`, `realvncconnect`, `synologyactiveprotectagent`, `tigervnc`, and `zoomvdiplugin`.

RETIRED INTEGRATION Removes 29 deprecated labels including `bartender`, `microsoftteams`, `nomad`, `nomadlogin`, `realvncviewer`, `skype`, and `vmwarehorizonclient`.

Renames the label file `npssantafull.sh` to `npssanta.sh`.

UPDATED INTEGRATION Updated application labels

└──▷ BREAKING ON UPGRADE

! The label file `npssantafull.sh` is renamed to `npssanta.sh`; any scripts or policies referencing `npssantafull` must be updated to `npssanta`.

! The following labels are deleted and can no longer be used: `audacity`, `bartender`, `basictex`, `code42`, `codex`, `debookee`, `eshareosx`, `facebookmessenger`, `goland`, `harmonysase`, `horos`, `island`, `mactex`, `magicbullet`, `maxcycling74`, `microsoftteams`, `nomad`, `nomadlogin`, `nps-santa`, `origin`, `realvncviewer`, `sketchupviewer`, `skype`, `steinbergcubaseelementsaile13`, `steinbergcubaseproartist13`, `techsmithcapture`, `vfx`, `vmwarehorizonclient`, `vuescanintel`.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ Deploy Apple declarations with Custom Declarations

NOTES ↗

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

Iru adds a Custom Declarations Library Item for deploying arbitrary Apple DDM JSON configurations, with real-time per-line schema validation.

NEW SUPPORTED PLATFORM Deploy any Apple DDM configuration before a dedicated Library Item exists by pasting or uploading declaration JSON into the new Custom Declarations Library Item.

Iru now validates each declaration against Apple's schemas as you type, flagging errors on the exact line.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Menlo Security

Feedback ⓘ announced · Nov 13, 2026 PRESS ↗

RANK

Secure your agentic enterprise with the Menlo Browser Security Platform. Get zero-day threat prevention, secure access, and data security in one unified platform.

> Menlo Security launched the IRIS Browser Extension, a centrally managed layer that gives security teams visibility and control over GenAI use, third-party extensions, and data movement across any browser without requiring a browser swap.

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

01 **In-session data-loss controls** NEW 65

Adds data-loss controls in the browser session, including clipboard, print, upload, and download restrictions, watermarking, read-only pages, developer tools restrictions, and screenshot deterrence.

— Names every control but not how to configure them launch-20261113-46aa324c

02 **IRIS Browser Extension deployment and cross-browser support** NEW 60

Menlo Security adds the IRIS Browser Extension, a centrally managed extension deployable via the Chrome Web Store, Microsoft Intune, JAMF, or other MDM/UEM tools, working across Chrome, Edge, Safari, Firefox, and Chromium-based browsers without replacing the user's browser.

— Names deployment channels and browsers but no further mechanism

launch-20261113-46aa324c

THINNER COVERAGE BELOW

03 **Inline file inspection with CDR** NEW 55

Adds inline file inspection for DLP violations, viruses, and malware, with optional Content Disarm and Reconstruction (CDR) to deliver clean files instead of blocking them.

— Names mechanism (CDR) but no setup steps launch-20261113-46aa324c

04 **Forced authentication and centralized audit logging** NEW

55

Adds forced authentication to the extension and centralized web logging to produce audit trails for SOC 2, HIPAA, and GDPR reviews and cyber insurance questionnaires.

— Names compliance frameworks but no log format or access path launch-20261113-46aa324c

05 **HEAT Shield AI threat prevention** NEW

50

Adds optional HEAT Shield AI for threat prevention against zero-day phishing, social engineering, and malvertising at click time.

— Names threat types but not detection mechanism launch-20261113-46aa324c

06 **Enterprise extension inventory and risk scoring** NEW

45

Adds an enterprise-wide inventory of third-party browser extensions with real-time risk scores and management controls, giving security teams visibility into installed extensions across all user devices.

— Describes capability but no scoring methodology or UI path launch-20261113-46aa324c

07 **Web filtering and GenAI guardrails** NEW

45

Adds web filtering and GenAI guardrails that steer users toward sanctioned tools and control what data is submitted to AI services without blocking them outright.

— States behaviour but no policy mechanism or config surface launch-20261113-46aa324c

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Passbolt

Feedback ⓘ updated · Nov 13, 2026 NOTES ↗ RANK

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

> Passbolt's latest release adds passkey-based MFA with a read-only Offline Mode, encrypted metadata with zero-knowledge mode, SCIM provisioning, and a generally available Safari extension, alongside new PIN code and notes resource types, custom fields, secret history, and in-app edition switching.

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

01 **Passkey MFA and Offline Mode support** NEW 55

Passbolt now supports phishing-resistant passkey MFA, extending QuickAccess and Offline Mode to PIN codes as well. The new Offline Mode (beta) gives users secure read-only access and autofill for selected credentials when the server is unreachable.

– Describes mechanism and scope but no setup steps snapshot-20261113

02 **SMTP OAuth for Exchange Online and PingOne SSO** NEW 55

Passbolt adds SMTP OAuth support for Microsoft Exchange Online to authenticate outbound email without static passwords, and adds a `PingOne` SSO integration for single sign-on authentication.

– Names both integrations but no config steps given snapshot-20261113

03 **Encrypted metadata with zero-knowledge mode** IMPROVED 50

Passbolt adds encrypted metadata with zero-knowledge mode for administrators, promoted to stable in `5.4.0`, and adds shared metadata key rotation.

– Names version and key rotation but no mechanism detail snapshot-20261113

04 **PIN code and notes resource types** NEW 45

Passbolt adds a dedicated PIN code resource type for storing PIN credentials and a standalone `notes` resource type.

– Names both new resource types but no usage path snapshot-20261113

05 **In-app edition switching** NEW 40

Passbolt adds in-app edition management so administrators can switch between Community and Pro editions directly from the application.

— Names a clear in-app admin action, no exact menu path snapshot-20261113

06 **SCIM provisioning for user management** IMPROVED 35

Passbolt adds SCIM provisioning (initially beta, now out of beta) for automated user management.

— Lifecycle noted but no endpoint or config detail snapshot-20261113

07 **Secret history for resources** NEW 35

Passbolt adds secret history, letting users view previously stored versions of a secret.

— Explains capability but no interface or retention detail snapshot-20261113

08 **Safari browser extension** NEW 30

Passbolt adds a Safari browser extension, now generally available.



— Bare announcement with no install or usage detail snapshot-20261113

09 **Custom fields for resources (beta)** NEW 25

Passbolt adds custom fields (beta) for resources.

— One-line beta announcement with no scope detail snapshot-20261113

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ Email Templates Editor UX Improvements (2026-11-12)

BLOG ↗ RANK

Secure users, AI agents, and more with Auth0, a production-ready authentication and authorization platform built to help you ship faster and drive sign-ups.

> Auth0 overhauled the Email Templates editor with a dedicated detail page, live preview, variable search, and an agent-prompt handoff for coding agents.

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

01 Email Templates editor and preview UX NEW 65

A new Emails > Templates list in the Auth0 dashboard shows all email templates with status, and selecting one opens a dedicated detail page with separate editor and live preview tabs showing the From line, Subject line, and rendered email preview.

When iterating on a branded email template, switch to the live preview tab to see exactly how the rendered email, From line, and Subject look before deploying.

> 📍 In the Auth0 dashboard, go to Emails › Templates, select the template you want to edit, then click the Preview tab to see the rendered output alongside From and Subject fields.

— Names exact navigation path and tabs Email Templates Editor UX Improvements …

02 Copy agent prompt for email templates NEW 65

The Email Templates editor adds a Copy agent prompt button that copies a pre-built prompt including the template's variables, current body, and deploy steps, for handoff to a coding agent.

When delegating email template authoring to a coding agent, use the Copy agent prompt button to give the agent full context — variables, body, and deploy steps — in one action.

> 📍 In the Auth0 dashboard, go to Emails › Templates, select a template, then click Copy agent prompt to copy a prompt containing the template's variables, current body, and deploy steps for use with a coding agent.

— Exact button and usage steps given Email Templates Editor UX Improvements …

THINNER COVERAGE BELOW

03 Template variables panel NEW 40

A new Variables panel in the Email Templates editor lets users search and copy template variables, with common variables grouped together for quick access.

— Describes feature but no exact field names Email Templates Editor UX Improvements …

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL?

Clerk Feedback ⓘ

View organizations and update users with the Clerk MCP server (2026-11-12)

BLOG ↗ RANK

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

> Clerk expanded its MCP server so agents can now inspect organization data and directly update user records.

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

01 Organization and membership listing in MCP server NEW

35

Clerk's MCP server now supports listing organizations and memberships from an agent, allowing agent-driven workflows to query org structure.

— Describes capability but no tool/endpoint names given

View organizations and update users wit…

02 User profile and metadata updates via MCP server NEW

35

Clerk's MCP server now supports updating user profiles and metadata from an agent.

— Clear capability, no specific command or schema named

View organizations and update users wit…

WAS THIS USEFUL?

Keycloak

Feedback ⓘ

Keycloak Client Libraries 26.0.13 released (2026-11-12) BLOG ↗ RANK

Keycloak - the open source identity and access management solution. Add single-sign-on and authentication to applications and secure services with minimum effort.

> Keycloak's client libraries added support for OpenJDK 25 in release 26.0.13.

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

01 **OpenJDK 25 support in client libraries** IMPROVED 35

Keycloak Client Libraries 26.0.13 now support `OpenJDK 25`, with automated tests updated to run against it by default.

— Names the JDK version but no further mechanism or migration detail

Keycloak Client Libraries 26.0.13 relea…

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Hydra, Ory Kratos, Ory Keto, Ory Polis, Ory Elements v26.4.16 released

NOTES ↗ RANK

Discover the latest features, improvements, and updates to Ory.

> Ory's v26.4.16 release adds granular control over outbound HTTP request limits in Ory Hydra.

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

01 **Per-feature outbound HTTP limits in Hydra** NEW 60

Ory Hydra now supports configuring outbound HTTP limits per feature using `default_http_client` for shared settings, with feature-specific overrides available for individual features.

— Names config key but lacks list of overridable features

Ory Hydra, Ory Kratos, Ory Keto, Ory Po…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ updated · Nov 13, 2026 NOTES ↗ RANK

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

> Varonis introduces Agent Manifest, a configuration layer for defining AI agent boundaries and stopping out-of-scope actions at runtime.

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

01 **Agent Manifest for AI agent boundaries** NEW 40

A new configuration layer lets you define an AI agent's purpose, expected behavior, and scope, so that out-of-scope actions are detected via runtime drift and stopped before they execute.

.
.
.

— Describes purpose and behavior but no config keys or UI path snapshot-20261113

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ changelog · Nov 13, 2026 CHANGELOG ↗ RANK

RegScale is a continuous compliance automation platform that helps organizations manage security controls, audit evidence, and regulatory requirements in real time.

> RegScale's latest release adds CLI tooling for importing Burp scans and copying risks between instances, lets Trivy/Grype pull scan files straight from S3 with incremental skip-and-reprocess behavior, and introduces new Tenable.io, Wiz, JCAM and config-validation settings — alongside four breaking changes to scanner integration flags, keys and mop-up sync behavior.

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

01 S3-sourced and incremental Trivy/Grype imports NEW

90

Trivy and Grype scan imports can now read files directly from S3 using the new `--s3-bucket` and `--s3-prefix` flags combined with an AWS profile or keys on the command line. Imports also skip already-imported scan files automatically, with a `--reprocess` flag to force re-import of everything; runs that skip files no longer close findings.

Pull Trivy scan results directly from an S3 bucket using a named AWS profile, so the import runs without local copies of the scan files.

```
$ regscale trivy import_scans --s3-bucket my-scan-results --s3-prefix trivy/2026-11/ --profile prod-readonly
```

— Flags and behavior named, runnable example given changelog-20261113-7d9f797f

02 **Burp scan import via CSV/XLSX** NEW 88

A new `regscale burp import_burp_csv` command imports Burp scan results from a CSV or XLSX file reformatted as one row per issue, with a `--dry-run` flag to preview the import before writing findings.

Preview a Burp scan import before writing any findings, to verify scope and issue count before committing.

```
$ regscale burp import_burp_csv --file burp_results.xlsx --dry-run
```

— Exact command and flag given, runnable `changelog-20261113-7d9f797f`

03 **Risk copy commands between instances** NEW 88

New `regscale risks export_risks` and `import_risks` commands copy risks and their attachments between RegScale instances; risks are matched by title or with `--risk_model_id`, ownership can be assigned with `--risk_owner_id`, and re-runs do not duplicate risks.

Move risks and their attachments from a staging instance to production, assigning a specific risk model and owner.

```
$ regscale risks import_risks --risk_model_id 42 --risk_owner_id 7
```

— Named commands and flags, runnable example `changelog-20261113-7d9f797f`

04 **ScannerIntegration close_outdated_findings flag renamed** 86

BREAKING

The `ScannerIntegration` flag for closing outdated findings is renamed from `close_outdated_findings` to `close_outdated_findings_enabled`; `close_outdated_findings` is no longer accepted as a flag name, and passing it to `set_keys` or `sync_findings` / `sync_assets` raises `TypeError`. A class-level or instance-level `close_outdated_findings = True` assignment now overwrites the inherited `close` method with a bool, and reading `close_outdated_findings` as a flag returns the method (always truthy) — both fail silently, so custom scanner integration code must be audited.

— Exact old/new flag names and failure modes given `changelog-20261113-7d9f797f`

05 **Validation of removed init.yaml keys** IMPROVED 76

`regscale config validate` now detects removed `init.yaml` keys — including `cisaKev`, `jiraDefaultIssueType`, and `jiraDefaultPriority` — reporting them as `removed_key` and logging a warning naming what to use instead.

— Command and exact keys named, no runnable example `changelog-20261113-7d9f797f`

---

06 **Databricks mop-up sync changes for CIS/CVE** BREAKING 75

In Databricks `sync_findings --mop-up`, a plan that already holds issues under the earlier combined source gets a second set of issues on its first run after upgrading, so operators must close or delete the old ones first. `sync_assets --mop-up` is now refused unless the export holds both CIS and CVE types.

— Exact command flags and migration step given `changelog-20261113-7d9f797f`

---

07 **defender export_resources -n flag meaning changed** BREAKING 73

On `defender export_resources`, the `-n` short flag now means `--query_name` only; scripts that used `-n` to skip the upload must switch to `--no_upload`.

— Exact flags and migration path named `changelog-20261113-7d9f797f`

---

08 **Tenable.io STIG compliance export mapping** NEW 70

DISA STIG audit results from the Tenable.io compliance export can now be sent to matching control objectives by enabling the new `tenableIoStigComplianceExport` setting, which is off by default.

— Names setting and default, lacks usage example `changelog-20261113-7d9f797f`

---

09 **JCAM @drop field mapping** IMPROVED 65

Mapping a JCAM value to `@drop` in `jcamValueMapping` now silently discards it, leaving the destination field unchanged instead of writing the literal text `@drop`.

— Config key and behavior named, no example `changelog-20261113-7d9f797f`

---

10 **Xray vulnerability keys now include CVE** BREAKING 63

Each Xray vulnerability is now stored under a key that includes its CVE; the first Xray import after upgrading creates new vulnerabilities with a new first-seen date and closes the old ones.

— Behavior change explained, no exact key format given `changelog-20261113-7d9f797f`

THINNER COVERAGE BELOW

---

11 **Wiz concurrent write quiet period** NEW 57

A new `wizHeldFindingsQuietSeconds` setting (default 30 seconds) controls how long Wiz waits for concurrent asset writes.

— Setting named with default but no usage example `changelog-20261113-7d9f797f`

12 **Settings tab in CLI GUI** NEW 53

The interactive terminal GUI gains a new Settings tab to browse and edit every CLI configuration key, grouped by integration, with secrets masked.

— UI path described but no exact navigation steps `changelog-20261113-7d9f797f`

└──▷ **BREAKING ON UPGRADE**

! The `ScannerIntegration` flag for closing outdated findings is renamed from `close_outdated_findings` to `close_outdated_findings_enabled`; `close_outdated_findings` is no longer accepted as a flag name, and passing it to `set_keys` or `sync_findings` / `sync_assets` raises `TypeError`. A class-level or instance-level `close_outdated_findings = True` assignment now overwrites the inherited `close` method with a bool, and reading `close_outdated_findings` as a flag returns the method (always truthy) — both fail silently, so audit all custom scanner integration code.

! On `defender export_resources`, the `-n` short flag now means `--query_name` only; scripts that used `-n` to skip the upload must switch to `--no_upload`.

! Each Xray vulnerability is now stored under a key that includes its CVE; the first Xray import after upgrading creates new vulnerabilities with a new first-seen date and closes the old ones.

! In Databricks `sync_findings --mop-up`, a plan that already holds issues under the earlier combined source gets a second set of issues on its first run after upgrading — close or delete the old ones first. `sync_assets --mop-up` is refused unless the export holds both CIS and CVE types.

WAS THIS USEFUL? 👍 👎

AGENT

◆ Agent Sandboxes

Alibaba Cloud OpenSandbox

Feedback ⓘ release-1.2.1

NOTES ↗ RANK

OpenSandbox is a join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> OpenSandbox's 1.2.1 release cycle adds per-request Credential Vault Exchange, multi-version Firecracker VMM bundles, and background-command polling in the Python MCP SDK, alongside stricter request validation, port-scoped egress enforcement, and controller leader-election tuning.

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

01 Controller leader-election tuning and extra args NEW

86

Controllers can now be tuned via `controller.extraArgs` (and matching Helm chart values) plus new flags `--leader-elect-lease-duration`, `--leader-elect-renew-deadline`, and `--leader-elect-retry-period`, with namespace-scoped caches also added. A 10-second API-server stall no longer forces a full re-election, and defaults are unchanged.

Pass extra controller flags at deploy time — for example to tune leader-election timings — without modifying the base chart values.

yaml

```yaml
controller:
  extraArgs:
    - --leader-elect-lease-duration=20s
    - --leader-elect-renew-deadline=15s
```

Tune leader-election on a slow control plane so a temporary API-server stall does not trigger unnecessary re-elections.

```
$ opensandbox-controller --leader-elect-lease-duration=20s --leader-elect-renew-deadline=15s --leader-elect-retry-period=5s
```

— Names exact flags, config key and a runnable command release-1.2.1-rc.2 release-1.2.1

02 Strict validation rejects unknown request fields

BREAKING

80

Unknown fields in closed request objects are now rejected with a 422 response instead of being silently discarded. Ten models — image, auth, platform, volume, OSSFS, credential-proxy, and snapshot-renew requests — now use `extra = "forbid"`, so options such as `readonly: true` on a volume and `createIfNotExist: false` behave as declared.

— Names response code, flag and exact models affected `release-1.2.1-rc.2` `release-1.2.1`

03 **Background command status polling in MCP SDK** NEW 75

The Python MCP server/SDK now exposes `command_get_status` (running state, exit code) and `command_get_logs` (content and opaque cursor) tools for polling long-running background commands and streaming incremental log output.

Poll the status and tail incremental logs of a long-running background command inside a sandbox using the Python MCP SDK.

```python
status = await sandbox.command_get_status(command_id)
logs = await sandbox.command_get_logs(command_id)
```

Poll a long-running background command for its exit code and stream incremental log output from the Python MCP server.

```python
# After launching a background command, poll status and logs
status = await mcp.call_tool('command_get_status', {'command_id':
'<cmd-id>'})
logs = await mcp.call_tool('command_get_logs', {'command_id':
'<cmd-id>', 'cursor': None})
```

— Named tools with working code example `release-1.2.1-rc.2` `release-1.2.1`

THINNER COVERAGE BELOW

04 **Multi-version Firecracker VMM bundles** NEW 55

Managed Firecracker 1.16.1 and 1.17.0 bundles let new template builds select the highest installed VMM version, while existing artifacts restore with their exact recorded VMM version, failing explicitly if that version is unavailable.

— Clear before/after behaviour but no runnable step `release-1.2.1`

05 **Docker writable-layer disk limit enforcement** BREAKING 55

Docker now enforces `resourceLimits.disk` on the writable container layer, returning 400 for unsupported drivers or invalid limits instead of silently ignoring them.

— Names config key and error behaviour, no example release-1.2.1

06 **Detailed download status and range APIs in SDKs** NEW

50

Detailed download APIs now expose status, headers, ranges, and partial-response detection in the Python, JavaScript, and C# SDKs without requiring changes to existing adapters; these SDKs report unknown totals when encoded lengths cannot describe decoded output.

— Names three SDKs and fields but no code sample release-1.2.1

08 **SHA-256-pinned OCI template artifacts required** BREAKING

30

OCI template artifacts must now be SHA-256-pinned in the expected repository.

— One-line requirement with no mechanism detail release-1.2.1

09 **Atomic staged uploads prevent overwrite data loss** IMPROVED

30

Staged uploads now reduce overwrite data loss where atomic replacement is available, protecting existing content on file replace.

— Thin description, no concrete surface named release-1.2.1

10 **Docker snapshots from paused sandboxes** NEW

20

Docker snapshots can now be taken directly from paused sandboxes.

— Bare one-line capability with no further detail release-1.2.1

└──▷ **BREAKING ON UPGRADE**

! Kotlin and C# SDK consumers must recompile and review widened credential types after the `CredentialSource` model was broadened to cover inline and exchange variants.

! C# SDK: inline-only `CredentialSource` assumptions must be updated when recompiling against the common model.

WAS THIS USEFUL? 👍 👎

microsandbox

Feedback ⓘ v0.7.15 NOTES ↗ RANK

> microsandbox v0.7.15 adds opt-in candidate installs and publishing, and reduces per-file host operations on passthrough mounts.

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

01 **Opt-in candidate installs and publishing** NEW 35

microsandbox now supports opt-in candidate installs and publishing, allowing pre-release versions to be installed and distributed through the existing release pipeline.

— Describes mechanism but no flag or command named v0.7.15

02 **Reduced per-file host operations on passthrough mounts**

IMPROVED 30

microsandbox now cuts per-file host operations on passthrough mounts, reducing filesystem overhead for workloads that use passthrough volume mounts.

— Performance claim without measurements or config detail v0.7.15

WAS THIS USEFUL? 👍 👎

Tencent Cloud CubeSandbox

Feedback ⓘ v0.8.0 NOTES ↗

RANK

> CubeSandbox v0.8.0 introduces a JuiceFS copy-on-write storage backend, filesystem-only snapshots, sandbox restart policies, a scheduler plugin system, and sandbox-to-sandbox TCP access via port mapping, alongside new E2B-compatible v2 API endpoints, per-sandbox QoS resource caps, and CubeEgress L7 rule improvements.

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

01 **E2B-compatible v2 sandbox endpoints** NEW 90

CubeSandbox now supports the E2B `/v2/sandboxes` and `/v2/sandboxes/{id}/connect` endpoints, enabling direct sandbox create and connect for E2B SDK 2.51.0 and later with no adapter layer.

Create and connect to a sandbox using the new E2B-compatible v2 endpoints without any adapter layer.

```
$ curl -X POST https://<cubesandbox-host>/v2/sandboxes \
  -H 'Content-Type: application/json' \
  -d '{"templateId": "<template-id>"}'

curl -X GET https://<cubesandbox-host>/v2/sandboxes/<sandbox-id>/connect
```

— Named endpoints with runnable curl examples v0.8.0

02 **Force flag for stuck template deletion** NEW 80

`cubemastercli tpl delete` gains a `--force` flag that removes template metadata and artifact copies on every node when a deletion task is stuck or in a bad build state.

Force-remove a template whose deletion task is stuck, clearing metadata and artifact copies cluster-wide.

```
$ cubemastercli tpl delete --force <template-name>
```

— Exact CLI flag and command given v0.8.0

03 **Private CA trust for Helm image pulls** NEW 70

The Helm chart adds a `trustedCACerts` setting so Kubernetes deployments can pull images from registries secured with private CAs.

Allow image pulls from an internal registry that uses a private CA in a Kubernetes Helm deployment.

```yaml
trustedCACerts: |
  -----BEGIN CERTIFICATE-----
  <your-CA-cert-PEM>
  -----END CERTIFICATE-----
```

— Named config key with values.yaml example v0.8.0

04 **Node overcommit configuration via ops-agent** NEW 70

A new resident `ops-agent` process receives commands from CubeOps and updates the Cubelet config file, allowing node overcommit to be configured from the CLI or Web UI.

— Names component and config file, no exact command v0.8.0

05 **Per-sandbox QoS resource caps** NEW 65

New QoS settings let template authors cap per-sandbox network bandwidth, packet rate, disk throughput, and IOPS at the template level, so a single sandbox cannot saturate shared host resources.

— Explains mechanism but no exact config key shown v0.8.0

06 **CubeEgress L7 rule enhancements** IMPROVED 60

CubeEgress L7 rules now support a `*` match-all wildcard, forward WebSocket upgrade headers, and can print the sandbox ID in audit logs.

— Names three concrete rule behaviours, no exact syntax v0.8.0

THINNER COVERAGE BELOW

07 **Run-as user for template sandboxes** NEW 50

A template can specify a run-as user, by UID or name, so sandboxes launched from that template run as that user.

— Clear behaviour but no config key shown v0.8.0

08 **Artifact distribution concurrency control** NEW 45

CubeMaster adds a setting to control artifact distribution concurrency, tuning how many nodes receive rootfs artifacts in parallel.

— No config key or default value named v0.8.0

09 **Snapshot list filtering and deletion by name** IMPROVED

40

Snapshot lists can be filtered by snapshot name or ID, and snapshots can be deleted by snapshot ID or name.

— Described only in prose, no exact surface v0.8.0

10 **Explicit storage backend selection for templates** NEW

35

Templates can now be created with an explicit storage backend specified via the SDK or API.

— No backend names or API shape given v0.8.0

WAS THIS USEFUL? 👍 👎

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 5.0.0 (2026-11-12) BLOG ↗ RANK

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

> Runlayer 5.0.0 ships lossless SIEM session export with payload URIs, custom TypeScript eval scorers over MCP, config-exception routing for admins, and always-on scoped agent token audience enforcement, alongside several breaking changes to Terraform and migration behavior.

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

01 Lossless SIEM session export via payload URIs

BREAKING

93

Export full session event bodies to `<prefix>/payloads/` via the new `data_payload_uri` field, so oversize and chunked events are no longer truncated in SIEM destinations. Before upgrading, SIEM export destinations must allow writes to all of `<prefix>/*` in the destination role or bucket policy, and consumers must filter S3-notification ingestion on `.jsonl.gz` (or `<prefix>/year=`) and dedupe on `session_id + event_id + part_index`.

Before upgrading, broaden your SIEM destination bucket policy to accept the new payload objects so full session event bodies export without permission errors.

```
$ # Allow writes to all objects under the export prefix, not just
*.jsonl.gz
# Then filter S3 notifications and dedupe on session_id +
event_id + part_index
aws s3api put-bucket-policy --bucket my-siem-bucket --policy '{
  "Statement": [{
    "Effect": "Allow",
    "Action": ["s3:PutObject"],
    "Resource": "arn:aws:s3:::my-siem-bucket/runlayer-sessions/*"
  }]
}'
```

— Names field, path pattern and exact migration requirements 5.0.0 (2026-11-12)

02 Scoped agent token audience enforcement always on

BREAKING

90

Scoped agent token audience enforcement is now always on: the `agent-token-audience-enforcement` flag is removed, `FF_DISABLE=agent-token-audience-enforcement` does nothing, and any scoped agent or on-behalf-of token used outside its scope is rejected with `403 invalid_target`. Clients must request a `resource` that matches every server they use the token on, or request a token without `resource`.

— Exact flag, error code and required client behavior named `5.0.0 (2026-11-12)`

---

03 **Custom agent configuration exception routing** NEW 88

Set a custom agent configuration exception URL with `agent_config_exception_request_url` on the Terraform `runlayer_organization` resource, or in Settings > General under Agent configuration exceptions in the UI, to route policy-exception requests to your own GitHub repo or form instead of in-app requests.

Route agent configuration exception requests to your own form or repo instead of handling them in-app — set once in Terraform and all policy-blocked settings link there automatically.

hcl

```hcl
resource "runlayer_organization" "main" {
  agent_config_exception_request_url = "https://github.com/your-org/ai-policy-exceptions/issues/new"
}
```

— Exact config key, Terraform resource and UI path given `5.0.0 (2026-11-12)`

---

04 **Devcontainer LLM Gateway URL pinning** BREAKING 88

Pin the devcontainer LLM Gateway to specific URLs by setting both `host` and `gatewayUrl` on the `llm-gateway` Feature, so developer device keys and prompts go only to those URLs. Setting only one of the two now fails the build.

Lock a devcontainer LLM Gateway to known-good URLs so developer keys and prompts never reach unexpected hosts.

json

```json
{
  "features": {
    "ghcr.io/runlayer/devcontainer-features/llm-gateway:1": {
      "host": "gateway.corp.example.com",
      "gatewayUrl": "https://gateway.corp.example.com/v1"
    }
  }
}
```

— Exact config fields and failure condition named `5.0.0 (2026-11-12)`

05 **Terraform `runlayer_server` field clearing change** BREAKING 83

Removing `icon_url` or `transport_config` from a `runlayer_server` Terraform config no longer clears those values; to clear them you must explicitly set `icon_url = ""` or `transport_config = {}`. This requires the next Terraform provider release.

— Exact resource, fields and remediation syntax given `5.0.0 (2026-11-12)`

06 **Sensitive data masked in audit trails and SIEM exports** IMPROVED 70

MCP URL credentials (sign-in details and tokens) are now hidden from audit-only viewers in Shadow AI listings and MCP discovery audit events, including previously recorded events. Agent environment variable values set through the `update_agent` Runlayer MCP tool no longer appear in tool-call audit rows or SIEM exports, with only variable names recorded.

— Names the MCP tool and affected surfaces but no config to act on `5.0.0 (2026-11-12)`

07 **GitHub CLI and TLA+ tools in agent sandboxes** NEW 60

Agent sandboxes now include the GitHub CLI (`gh`) and TLA+ tools (`tlc` , `sany` , `pcal`), so agents can script GitHub workflows and model-check specs without installing anything.

— Names tools but gives no usage example `5.0.0 (2026-11-12)`

THINNER COVERAGE BELOW

08 **Plugin table migration locks and blocks rollback** BREAKING 55

The database migration in this release briefly locks the plugin table and does not support rollback; an older version cannot load the Runlayer Plugin after the migration runs. Roll forward to fix problems, or contact Runlayer support before attempting a rollback.

— Describes impact but no timing, duration or table name `5.0.0 (2026-11-12)`

09 **Offline MCP call decisions reported on reconnect** NEW 55

Offline MCP call decisions made by AI Watch on a device are now reported once the device reconnects, producing an audit log entry marked Decided on the device .

— Explains mechanism and audit marker, no navigation given `5.0.0 (2026-11-12)`

10 **Custom TypeScript eval scorers via MCP** NEW 50

Create, delete, and draft custom TypeScript eval scorers through the Runlayer Assistant or any MCP client using the Runlayer MCP server.

— Names surface and verbs but no example or schema `5.0.0 (2026-11-12)`

11 **Concurrent connector onboarding in Go** IMPROVED 45

Connect multiple connectors in Go concurrently, with each connector card progressing from Add to Adding to Added independently and back-to-back additions batched into one reply.

— Describes UI behavior but no exact navigation path `5.0.0 (2026-11-12)`

12 **Failed tool calls included in session access history** IMPROVED 35

Session-based access policies now account for tool calls that reached a connector and then failed, so failed calls are included in session history.

— States behavior change only, no scope or mechanism detail `5.0.0 (2026-11-12)`

13 **Faster LLM routing rule propagation** IMPROVED 35

LLM routing rule changes now take effect the next time a device requests a gateway token, rather than waiting for its next check-in.

— Simple timing change, no config or numbers given `5.0.0 (2026-11-12)`

└──▷ BREAKING ON UPGRADE

! Scoped agent token audience enforcement is now always on: the `agent-token-audience-enforcement` flag is removed, `FF_DISABLE=agent-token-audience-enforcement` does nothing, and any scoped agent or on-behalf-of token used outside its scope is rejected with `403 invalid_target`. Clients must request a `resource` that matches every server they use the token on, or request a token without `resource`.

! SIEM export destinations must allow writes to all of `<prefix>/*` in the destination role or bucket policy before upgrading, and consumers must filter S3-notification ingestion on `.jsonl.gz` (or `<prefix>/year=`) and dedupe on `session_id + event_id + part_index`.

! Setting only `host` or only `gatewayUrl` on the `llm-gateway` Feature now fails the build; both must be set together.

! Removing `icon_url` or `transport_config` from a `runlayer_server` Terraform config no longer clears those values; to clear them you must explicitly set `icon_url = ""` or `transport_config = {}`. Requires the next Terraform provider release.

! The database migration in this release briefly locks the plugin table and does not support rollback; an older version cannot load the Runlayer Plugin after the migration runs. Roll forward to fix problems, or contact Runlayer support before attempting a rollback.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

Straiker

Feedback ⓘ

Straiker EU & India Expansion | Regional AWS Hosting (2026-11-12) BLOG ↗

RANK

Straiker discovers and inventories enterprise AI agents, tests them with automated red teaming and protects them with runtime guardrails.

> Straiker expanded its regional hosting footprint with new AWS Frankfurt and upcoming AWS Mumbai regions, paired with customer-managed encryption keys for data sovereignty.

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

01 **EU and India regional AWS hosting** NEW 50

Straiker now hosts in AWS Frankfurt (available from November 2026) to support EU data residency requirements, and is adding an upcoming AWS Mumbai region to give Indian organizations a local hosting option for agentic AI security.

— Names regions and timing but no setup steps Straiker EU & India Expansion | Regiona…

02 **Customer-managed encryption keys** NEW 40

Straiker adds customer-managed encryption keys so organizations retain control over their own cryptographic key infrastructure and can prevent Straiker from accessing or decrypting their data without authorization.

— Describes capability but no config or key management details

Straiker EU & India Expansion | Regiona…

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal

Feedback ⓘ

release-049c4fbd · NOTES ↗ · RANK

Minimal is a secure, declarative runtimes for Claude Code, Gemini, Codex, Pi and more, so you can run multiple agents safely on your stack from a verified software supply chain.

> Minimal shipped box audit trail inspection, GitHub-backed authentication, automatic update checks, redacted bug reporting, and explicit network egress controls for task runs.

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

01 **Egress control for task runs** NEW 80

`min task run` supports `--network none` and `--egress` flags to restrict outbound traffic during a task, so only explicitly allowed destinations (e.g. `--egress github`) can leave the box.

Run a task with explicit egress control so only allowed outbound traffic leaves the box.

```
$ min task run --network none --egress github mybuildtask
```

— flags and example command given, broader scope undocumented release-049c4fbd

02 **Audit trail reading with `min box audit`** NEW 75

`min box audit` reads a box's audit trail, and `min box audit --parent` merges a box's parent chain so you can inspect the full audit history, including ancestor boxes, to verify no records were edited.

Inspect the full audit trail for a box, merging its parent chain, to verify no records were edited.

```
$ min box audit --parent mybox
```

— exact command given, record format/mechanism unexplained release-049c4fbd

03 **Bug reports bundle telemetry spool** NEW 60

`min bug` carries the telemetry spool alongside the report, with telemetry settings redacted automatically before sending.

— command named, redaction behaviour stated but no example release-049c4fbd

WAS THIS USEFUL?  

◆ AI Red Teaming & Testing

humanbound

Feedback ⓘ v2.15.0 NOTES ↗ RANK

Humanbound is an open-source adversarial testing engine, SDK, and CLI for AI agents. Runs locally or against the Humanbound Platform.

> Humanbound's v2.15.0 release adds a full CLI workflow for building and validating custom attack orchestrators, richer subscription introspection, and a new orchestrator alias for behavioral QA testing.

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

01 Custom orchestrator management via `hb orchestrators` NEW

90

A new `hb orchestrators` command family lets users manage custom attack orchestrators end-to-end: `hb orchestrators list` shows built-in and custom orchestrators including their version and description from `manifest.yaml` in `~/.humanbound/orchestrators/`; `hb orchestrators new <name>` scaffolds a working single-turn attack defined in `config.py` and scored by an LLM judge; and `hb orchestrators validate <name|path>` checks a custom orchestrator's entry points, arguments and manifest without calling an LLM or agent, exiting 1 on failure so it can run in CI.

Check whether a custom orchestrator is valid before merging a PR — fails CI if any entry point, argument or manifest field is wrong.

```
$ hb orchestrators validate my-orchestrator
```

— Names exact commands, config paths and CI exit behavior. v2.15.0

02 Subscription status in `hb orgs subscription` IMPROVED

67

`hb orgs subscription` now shows subscription status (`active` , `past_due` , `unpaid` or `canceled`), pending plan switches with their effective dates, and per-feature values such as `sso: oauth` and `forensic_window_days: 90` .

— Names command and exact field values but no workflow detail. v2.15.0

THINNER COVERAGE BELOW

03 Behavioral QA orchestrator alias NEW

42

The Behavioral QA orchestrator can now be selected with `-t behavioral_qa` , in addition to the already-supported `-t qa` and `-t behavioral` aliases.

— Names the flag but no further mechanism given. v2.15.0

WAS THIS USEFUL? 👍 👎

◆ AI Security Posture & Discovery

Onyx Security

Feedback ⓘ announced · Nov 12, 2026 PRESS ↗

RANK

Onyx Security is a control plane that discovers, monitors and governs enterprise AI agents inline across browsers, coding assistants, cloud workloads and MCP servers.

> Onyx Security extended its Secure AI Control Plane with dedicated support for Snowflake Cortex, adding both inventory/monitoring of Cortex Agents and runtime enforcement for the Cortex Code CLI.

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

01 **Cortex Code CLI runtime enforcement** NEW 67

A new Snowflake Cortex Code CLI integration automatically detects Cortex Code CLI activity and enforces runtime security controls: inspecting prompts against enterprise policies, blocking policy-violating tool calls before execution, and applying MCP server sanctioning policies to govern approved and unapproved MCP server interactions.

— Enforcement mechanism described but no CLI flag or config key given

launch-20261112-7abf297f

02 **Snowflake Cortex Agents discovery and monitoring** NEW 65

A new integration automatically inventories Cortex Agents across a Snowflake account, mapping their connected tools, MCP servers, Cortex Search and Analyst data sources, and user access permissions. It also captures complete interaction transcripts — prompts, responses, tool invocations, and SQL calls — and evaluates them against security policies to generate alerts for unsafe or unauthorized behavior, with this visibility surfaced in the Secure AI Control Plane alongside other agent environments.

— Mechanism and data sources named, but no config or command shown

launch-20261112-7abf297f

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

GitHub CLI

Feedback ⓘ v2.104.0 NOTES ↗ RANK

GitHub CLI github's official command line tool

> GitHub CLI v2.104.0 adds a safety flag to `gh repo delete` that guards against accidental deletion of private forks, which also changes behavior for existing non-interactive scripts.

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

01 Cascading deletion confirmation flag for repo delete

BREAKING

88

`gh repo delete` gains a new `--accept-cascading-deletion-of-private-forks` flag that must be passed alongside `--yes` to confirm cascade warnings when deleting a private or internal repository that has private forks, enabling non-interactive and agent workflows. This is a breaking change: existing scripts that pass only `--yes` will now stall waiting for confirmation unless the new flag is added.

Delete a private repo that has private forks non-interactively in a CI or agent workflow without manual prompts.

```
$ gh repo delete OWNER/REPO --yes --accept-cascading-deletion-of-private-forks
```

— Names exact flag, command, and breaking impact with runnable example. v2.104.0

└──▷ **BREAKING ON UPGRADE**

! `gh repo delete --yes` on a private or internal repository with forks now requires the additional `--accept-cascading-deletion-of-private-forks` flag to proceed non-interactively; existing scripts that pass only `--yes` will stall waiting for confirmation.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

WAF - Manage AI Security for Apps with Terraform (2026-11-12) BLOG ↗ RANK

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

> Cloudflare shipped Terraform support for AI Security for Apps, expanded WAF malicious uploads detection, and added new D1, Rules, and Kitesurf capabilities across its developer platform.

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

01 **AI Security for Apps managed via Terraform** NEW 95

Terraform can now manage AI Security for Apps end to end: the `cloudflare_ai_security` resource enables or disables it per zone (Cloudflare provider v5.29.0), the `cloudflare_ai_security_custom_topics` resource defines custom topics each prompt is scored against using `label` and `topic` fields, and AI Security for Apps fields inside `cloudflare_ruleset` let detection settings, topics, and mitigation rules live in one configuration.

Turn on AI Security for Apps for a zone and define a custom competitor-detection topic in a single Terraform configuration.

hcl

```
resource "cloudflare_ai_security" "zone_ai_security" {
  zone_id = var.cloudflare_zone_id
  enabled = true
}

resource "cloudflare_ai_security_custom_topics"
"zone_custom_topics" {
  zone_id = var.cloudflare_zone_id
  topics = [
    {
      label = "competitors"
      topic = "seeking info on competitors"
    }
  ]
}
```

— Names exact resources, fields, provider version, with runnable example

WAF - Manage AI Security for Apps with …

02 **Malicious uploads detection flags boundary-bearing non-multipart headers**

IMPROVED — 85

WAF malicious uploads detection now sets `cf.waf.content_scan.has_failed` to `true` for requests whose Content-Type header includes a `boundary` parameter but whose type is not `multipart/*` (for example, `application/json; boundary=example`), marking the request body as unscanned.

— Exact field name and trigger condition given, no runnable example

WAF - Malicious uploads detection flags…

THINNER COVERAGE BELOW

03 **Hash functions in Rules expressions** NEW — 55

Rules expressions now support sha1(), md5(), and sha512() functions to compute hashes directly within rule expressions.

— Named functions are directly usable, no example expression given product docs

04 **D1 export of databases with virtual tables** NEW — 50

D1 databases containing virtual tables such as FTS5 can now be exported as a SQLite CLI-compatible dump.

— Names FTS5 and SQLite CLI format but no export command shown product docs

05 **Kitesurf browser engine open-sourced** NEW — 50

Kitesurf, Cloudflare's experimental browser engine for AI agents, is now open source and can be run locally or deployed to a Cloudflare account.

— Names deployment options but no repo link or command given product docs

06 **D1 database as AI Search data source** NEW — 45

A Cloudflare D1 database can now be indexed in AI Search (beta) using a read-only SQL query as the data source.

— Describes mechanism but no config or command shown product docs

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue