THE CYBER TOOLCHAIN

NO. 135

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 10, 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 168 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 39 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. Socket now builds SBOMs from Gradle, sbt, Maven and .NET build output instead of manifests, so fixes

vim ? help

and reachability reflect what actually compiled. Lakera's second-stage model reads your system prompt before flagging a prompt attack, cutting false positives 95% on average. Yubico's Windows background service drops the admin elevation FIDO2 used to demand.

IDENTIFY

Get a dependency inventory from what the build actually produced

Socket infers SBOMs per build root from Gradle, sbt, Maven and .NET build output, closing the long-standing gap where a manifest lists a version range and the resolved graph contains something else. ZeroPath pairs reachability-aware scanning across 13 ecosystems with AI-BOM and container image inventory, so the list of what you run includes the model and image layers, not just the lockfile.

Socket · ZeroPath

PROTECT

Stop babysitting prompt-attack alerts and FIDO2 elevation prompts

Lakera's two-stage Prompt Defense passes borderline input to a larger model with the application's own system prompt as context, which is the only way to tell a jailbreak from a user who legitimately asks the app to ignore instructions — 95% fewer false positives on average means the detector can stay in blocking mode. Yubico Authenticator 7.5.0's Windows background service removes the admin elevation FIDO2 operations required, which was the practical reason key management stalled on locked-down laptops; PIV also gains ML-DSA/ML-KEM.

Lakera · Yubico

PROTECT

Issue short-lived identity to agents and sessions instead of handing out keys

JumpCloud's Agent Delegated Access mints signed JWTs so an AI agent acts as a scoped identity you can revoke, rather than inheriting a human's long-lived token. P0 Security binds OAuth, native and impersonation credentials to a client owner and session, so a leaked credential is useless outside the session that minted it, and 1Password Environments gives a project a shareable secret collection reachable from CLI, SDK and MCP server instead of a shared vault.

JumpCloud · P0 Security

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

Feedback ⓘ updated · Nov 10, 2026 NOTES ↗

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

Hacktron adds a full pentest and PR-review API, an MCP tool suite, scheduled whitebox scans, bidirectional Jira sync, attestation letters, and finding IDs across all source control platforms.

Run a pentest end-to-end through the API with `POST /pentests`, `PATCH /pentests/{id}`, `POST /pentests/{id}/login-checks`, `POST /pentests/{id}/estimate`, and `POST /pentests/{id}/start`.

List open pull requests and trigger a Hacktron review programmatically with `GET /pull-requests`, `GET /pull-requests/details`, and `POST /pull-requests/review`.

Read and change PR review settings, including `scm_check_statuses_enabled`, and toggle reviews per repository with `GET /pr-review-settings`, `PATCH /pr-review-settings`, and `POST /repos/pr-reviews`.

Stop a scan, upload a code archive, discard a draft, and fetch credit and finding totals with `POST /scans/{id}/stop`, `POST /archives/uploads`, `DELETE /pentests/{id}`, `GET /credits`, and `GET /findings/counts`.

Start a fix pull request for a finding, track its progress, and check automation credits with `POST /findings/{id}/remediate`, `GET /findings/{id}/remediation`, and `GET /credits/automation`.

+21 more

NEW VULNERABILITY CHECK Transitive dependency CVE/advisory reachability checking in Slack

NEW SUPPORTED PLATFORM Dart, Haskell, Erlang syntax highlighting in findings and reports

NEW SUPPORTED PLATFORM Chinese, Japanese, and Korean text rendering in PDF reports

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Code Security

Snyk

Feedback ⓘ changelog · Nov 10, 2026 DOCS ↗ BLOG ↗ RANK

The official changelog for the Snyk AI Trust Platform. Discover what's new in our security intelligence, AI-powered tools, and latest product enhancements.

> Snyk shipped a Customer Gateway for scanning internal applications behind the network perimeter, alongside risk-based MCP Governance policies, a new MCP Risk Score, finer tenant role scoping, and signed installers for the ADS agent.

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

01 **Signed ADS installers for macOS and Windows** NEW 90

The ADS installer now ships as a signed macOS `.pkg` and Windows `.msi`, downloadable from Settings in Evo, and installable via `sudo installer -pkg <PACKAGE_FILE> -target /` or `Start-Process msiexec.exe` with `TENANTID` and `PUSHKEY` arguments.

Install the ADS agent on a macOS machine using the signed package, supplying your tenant credentials so it registers automatically.

```
$ sudo installer -pkg snyk-ads-installer-macos.pkg -target /
/usr/local/bin/snyk-ads/snyk-ads-installer \
  --tenant-id "<TENANT_ID>" \
  --push-key "<PUSH_KEY>"
```

Install the ADS agent silently on Windows with a verbose install log for auditing.

```
$ Start-Process msiexec.exe -Wait -ArgumentList '/i snyk-ads-
installer-windows-x86_64.msi TENANTID=<TENANT_ID> PUSHKEY=
<PUSH_KEY> /qn /l*v ads-install.log'
```

— Exact install commands and arguments make it directly runnable 2026-11-10-29977eee

02 **Customer Gateway for internal application scanning** NEW 75

Create a Customer Gateway from the CLI or from Settings and run it in your environment with Docker, Kubernetes, or the CLI to let Snyk Continuous Offensive Security reach applications only accessible inside your network. Greybox assessment is available through the gateway when repositories are linked, and every authentication method works through the gateway at the same depth as public targets.

— Names deployment methods and auth parity but no exact command 2026-11-10-29977eee

03 **Risk-based policies in MCP Governance** NEW 70

MCP Governance now supports risk-based policies, letting you classify servers with an allowlist, define conditions on Evo risk dimensions such as prompt injection in a tool or unofficial publisher status, and preview blocking impact before enabling it.

— Describes mechanism but only via UI, no config surface named 2026-11-10-29977eee

04 **MCP Risk Score in Agent Supply Chain Security** NEW 60

Agent Supply Chain Security now surfaces an MCP Risk Score combining server behavior, trust, and local-server implementation signals, usable as a policy condition alongside any individual risk index.

— Explains composition but not scoring mechanics or usage steps 2026-11-10-29977eee

THINNER COVERAGE BELOW

05 **Enable and disable Evo policies** NEW 55

Evo by Snyk now lets you disable and enable any policy, including default policies; a disabled policy stops raising new issues and hides previously raised issues until re-enabled.

— Clear behaviour described, no named flag or endpoint 2026-11-10-29977eee

06 **Organization-scoped Tenant roles for Evo** IMPROVED 55

The Tenant Member role now has read-only Evo access scoped to the user's Organizations, and both Tenant Member and Tenant Viewer roles can be scoped by Organization.

— Names roles and scope but no admin UI path given 2026-11-10-29977eee

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ Claude Fable 5.2 in GitHub Copilot (2026-11-09)

BLOG ↗ RANK

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

> GitHub added Claude Fable 5.2 model support to Copilot across its surfaces and extended secret scanning to catch leaked secrets in GitHub Actions workflow logs.

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

01 Secret scanning for Actions workflow logs NEW 85

Secret scanning now scans GitHub Actions workflow logs for provider patterns when a workflow run completes, in addition to issues, wikis, discussions, and pull requests. Findings appear in the repository's secret scanning alert list alongside code alerts, with a link back to the specific workflow log; validity checks run for supported secret types, and alerts remain available even after the workflow logs are deleted. Enable it via enterprise settings under `Policies > Advanced Security`, setting the All repositories dropdown to Enabled.

Enable Actions workflow log scanning for all repositories across your enterprise to catch secrets leaked into CI output.

> 📍 In the enterprise settings, go to Policies › Advanced Security, then set the All repositories dropdown to Enabled under Secret scanning for Actions workflow logs.

— Names exact admin path, behaviour and persistence details

Secret scanning detects secrets in GitH…

02 Claude Fable 5.2 in Copilot NEW 83

GitHub Copilot now supports Anthropic's Claude Fable 5.2 model for complex, multistep coding tasks across large codebases, available in Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, the Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse via the model picker. Copilot Enterprise and Business admins must enable the Claude Fable 5.2 policy in Copilot settings (off by default); eligible enterprise customers can get zero data retention via a time-bound exemption while Anthropic rolls out Enterprise Frontier Safeguards (EFS). It's available to Copilot Pro+, Max, Business, and Enterprise users, billed at provider list pricing under usage-based billing.

Enable Claude Fable 5.2 for your organization so members can select it in the model picker across supported surfaces.

📍 In GitHub Copilot settings (Enterprise or Business admin), navigate to Policies, find the Claude Fable 5.2 policy, and toggle it on. Note: enabling constitutes acknowledgment of Anthropic's data retention requirement.

— Lists surfaces, tiers, pricing and admin toggle path

Claude Fable 5.2 in GitHub Copilot (202…

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

fcaptcha

Feedback ⓘ v1.47.3 NOTES ↗ RANK

Fcaptcha is a self-hosted, invisible CAPTCHA that detects AI agents driving real browsers. 2 AA, and EAA: no cookies, no third-party data sharing, and no visual or audio puzzle.

> fcaptcha's latest release adds two research-oriented capabilities: bounded stdin imports for ephemeral exports and a research-only replay mode with per-sample isolation.

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

01 **Research-only replay mode with per-sample isolation** NEW 39

fcaptcha adds a research-only replay mode with per-sample isolation and no gates or persistent JSON, intended for research use rather than production.

— Names mechanism but no command or config surface given v1.47.3

02 **Bounded stdin imports for research exports** NEW 23

fcaptcha now accepts bounded stdin imports for ephemeral research exports.

— Describes capability only in prose, no flag or limit named v1.47.3

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ

Google Cloud Load Balancer module 1.9.0 (2026-11-09) BLOG ↗ RANK

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

> DataDome's Google Cloud Load Balancer module can now read the request host from the X-Forwarded-Host header instead of the default host.

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

01 **X-Forwarded-Host support in GCLB module** NEW 75

The Google Cloud Load Balancer module can read the request host from the `X-Forwarded-Host` header instead of the default host by setting `DATADOME_USE_X_FORWARDED_HOST` in the module configuration. This is useful when the load balancer sits behind a reverse proxy that sets X-Forwarded-Host, so DataDome evaluates the correct original hostname.

When your Google Cloud Load Balancer sits behind a reverse proxy that sets X-Forwarded-Host, enable this setting so DataDome evaluates the correct original hostname rather than the proxy's host.

env

```
DATADOME_USE_X_FORWARDED_HOST=true
```

— Names exact config variable and use case, but no deeper mechanism detail

Google Cloud Load Balancer module 1.9.0…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.7 NOTES ↗ RANK

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

> Infisical's v0.167.7 release adds native honey token management alongside new integrations for ClickHouse, OVHcloud, and Databricks, plus PKI and Agent Vault improvements.

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

01 **ClickHouse rotation, OVHcloud auth, Databricks gateway support**

NEW 65

Adds `ClickHouse` credentials rotation to the secret rotation engine, token authentication for OVHcloud KMS connections, and gateway routing support for Databricks app connections and syncs.

— Three named integrations but no setup steps given v0.167.7

02 **Agent Vault multi-project and access management** IMPROVED

65

Agent Vault now supports multiple projects per organization, exposes inline access management in the Manage Access sheet, shows inline services on sessions, and lets each Agent Vault member sign in independently to MCP OAuth services.

— Names UI surface but no API or config detail v0.167.7

THINNER COVERAGE BELOW

03 **Native honey token management** NEW 50

Infisical now manages honey tokens natively, letting teams deploy and revoke AWS IAM-based canary credentials without external tooling.

— Describes mechanism but no UI path or command given v0.167.7

04 **Cross-signed internal CA certificates** NEW 40

Support cross-signed certificates for internal CAs, enabling trust chain flexibility across PKI hierarchies.

— Describes benefit but no config or command shown v0.167.7

05 **DigiCert certificate reissue on same order** IMPROVED

40

Reissue DigiCert TLS certificates on the same order, avoiding new order creation on renewal.

— Clear behavior change but no command shown v0.167.7

06 **Collapsible recently deleted projects list** IMPROVED

40

Recently deleted projects now appear as a collapsible group in the project list.

— Clear UI location but minimal functional detail v0.167.7

07 **Zstd compression for secret cache** IMPROVED

35

Compress secret cache payloads with zstd to reduce cache size.

— States mechanism but no measurable impact given v0.167.7

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

juice-shop

Feedback ⓘ v20.4.0 NOTES ↗ RANK

OWASP Juice Shop: Probably the most modern and sophisticated insecure web application

> Juice Shop v20.4.0 introduces Passkey/WebAuthn passwordless authentication alongside two new passkey-themed hacking challenges, and overhauls Score Board filtering with technology tags and grouped categories.

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

01 **Technology tags and grouped Score Board filters** IMPROVED 68

Tags challenges with new Technology filter tags including `SQL`, `NoSQL/MongoDB`, `JWT`, `OAuth`, `2FA/TOTP`, and `XML`, and organizes all Score Board filter tags into named groups — Constraints, Flavor, Learning, Technique, and Technology — with translatable labels and delayed tooltips. Tutorial and Coding Challenge buttons on challenge cards are now highlighted when filtering by their corresponding tag.

– Names tags and groups, UI path implied but not exact v20.4.0

02 **Passkey/WebAuthn passwordless authentication** NEW 60

Adds Passkey/WebAuthn passwordless authentication with a dedicated passkey management screen that displays the provider name and icon, plus a new `passkey-user` customer account for use with passkey challenges.

– Names screen and account but no setup steps v20.4.0

THINNER COVERAGE BELOW

03 **Passkey Bypass and Passkey Hijack challenges** NEW 30

Adds a Passkey Bypass five-star challenge and a Passkey Hijack five-star challenge, each with an accompanying coding challenge.

– Named challenges but no mechanism described v20.4.0

04 **Edit button for AI chat messages** NEW 25

Adds a button to edit user messages in the AI chat.

– Bare description with no mechanism v20.4.0

WAS THIS USEFUL? 👍 👎

Dalfox

Feedback ⓘ v3.3.0 NOTES ↗ CODE ↗

Dalfox is an XSS scanner that automates parameter analysis, reflected/stored/DOM-based XSS detection, and WAF fingerprinting for web apps.

Dalfox adds OpenAPI/Postman spec scanning, actionable confidence filtering with `--min-confidence`, blind OOB and session-loss parity on server/MCP, per-parameter filter fingerprints in all output formats, repeatable `--sxss-url`, and `SCAN_TIMEOUT` error codes.

└──▷ GET THIS VERSION ▾

Homebrew

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

Docker

```
$ docker pull ghcr.io/hahwul/dalfox:v3.3.0
```

git

```
$ git clone --branch v3.3.0 https://github.com/hahwul/dalfox.git
# already have the repo? check out this version:
$ git checkout v3.3.0
```

Homebrew lists this version · published under the project's GitHub account · checked 10 Nov

Wrong command?

NEW INTEGRATION Scan OpenAPI 3.x / Swagger 2.0 specs or Postman collections with `-i openapi` or `-i postman`, and retarget the server URL with `--base-url`.

Filter findings by confidence across CLI, REST, MCP, `--limit`, and exit code with `--min-confidence <low|high>`; `meta.min_confidence.dropped` prevents filtered targets from appearing clean.

Server and MCP now support blind OOB (`blind_oob`, `blind_oob_wait`) and session-loss detection (`session_check`, `session_check_url`), matching CLI behavior.

Findings now include a per-parameter `filter` fingerprint (`allowed`, `blocked`, `escaped` characters) in every output format.

`--sxss-url` is now repeatable (config key `sxss_url` takes a list); file-upload injections now target both file content and filename.

+2 more

NEW DETECTION CSTI detection on AngularJS / Vue escaped reflections

NEW DETECTION JS exec-sink detection for eval / setTimeout / Function / javascript: navigation

UPDATED DETECTION DOM-XSS AST sink coverage expanded

UPDATED DETECTION Outdated library detection expanded to DOMPurify, Underscore, Knockout, jQuery Validation, Mustache.js

NEW DETECTION WAF fingerprints for DDoS-Guard and Reblaze; --force-waf adds Wallarm, NAXSI, SafeLine

---

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-11-07 CODE ↗

CAPEv2 is a malware Configuration And Payload Extraction

CAPEv2 adds a `max_executables` cap for archive analysis packages and configurable static preview with early dynamic analysis stop in the web UI

└──▷ **USE IT**

Cap how many executables CAPE extracts and runs from a submitted archive to bound analysis time and resource use.

```ini
max_executables=5
```

Prevent duplicate auto-mode processor supervisors from racing on the same completed-task queue — the lock file is created automatically; verify it exists after starting the processor in auto mode.

```
$ python utils/process.py auto
```

Limit the number of executables processed from an archive submission with the new `max_executables` option in the archive analysis package.

Configure static preview behaviour and enable early dynamic analysis stop from the web UI via the new configurable static preview feature (`feat(web): configurable static preview and early dynamic analysis stop`).

CAPEv2 now bundles `data/7zz` (7-Zip 26.06 x86_64, mode `0755`) directly in the core repository so modules such as `file_extra_info` , `msi_extract` , `suricata` , and web zip downloads work without fetching from `community` .

CAPEv2 now preloads `libmagic` handles once at module import time in `lib/cuckoo/common/path_utils.py` via _preload_magic(), and adds shared magic_from_buffer() and magic_from_file() helpers with an `os.memfd_create` fallback for ZIP and other EOF-indexed formats misidentified by `file` 5.46+.

Add `pcodedmp` integration to Office static analysis: when `pcodedmp` is installed, the `parse_office` module now uses it for P-code decompilation alongside the existing `oletools` stack.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform Feedback ⓘ

November 4, 2026 - Large list loading, product totals, and stripped-build components (20

BLOG ↗

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

Finite State Platform now documents filter and sort fields on list endpoints, returns stable 400 error codes for invalid filters, and flags estimated Audit Trail totals via a response header.

Filter and sort fields, operators, and forms on list endpoints are now documented; an invalid filter or sort returns HTTP 400 with the valid fields and a stable code — `FILTER_FIELD_UNKNOWN`, `FILTER_OPERATOR_INVALID`, `FILTER_INVALID`, `SORT_FIELD_UNKNOWN`, or `SORT_INVALID` — so callers can match on the code when message text differs.

On `GET /public/v0/audit`, `GET /public/v0/audit-events`, and `GET /public/v0/projects/{projectId}/audit`, a total above 10,000 is flagged as an estimate via the `X-Total-Count-Estimated: true` response header, and the Audit Trail now shows totals above 10,000 as `10,000+`.

NEW DETECTION Stripped-build component identification for TweetNaCl, mDNSResponder, ebisu, Mbed TLS, and legacy cJSON

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.16.3 NOTES ↗ CODE ↗ RANK

Netcap is a framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap's latest releases add a full `net firewall` WAF command with policy-driven audit responses, new `net split` and `net mcp` commands, TLS 1.2/1.3 decryption via embedded NSS key logs, and a string of capture-reliability, detection, and database-management improvements.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 Audit-record response policies in net capture NEW

83

Apply explicit audit-record response policies during capture with `--firewall-policy` and dispatch live response effects over `-iface` captures with `--firewall-effects-config` on `net capture`. Print the versioned audit policy field types, sensitivity, and projection bounds as JSON using `--firewall-schema`.

Inspect the firewall's audit policy schema before writing a policy file, to see which fields, sensitivity levels, and projection bounds are available.

```
$ net capture --firewall-schema
```

Apply a response policy to a completed capture and dispatch live effects on an active interface capture simultaneously.

```
$ net capture -iface eth0 --firewall-policy policy.json --
firewall-effects-config effects.json
```

— Three flags named with runnable examples for each. v0.16.1

02 New `net firewall` WAF command NEW

78

New `net firewall` command adds a WAF with routed HTTP/HTTPS and API inspection, scoped identities and approvals, a Coraza/HTTP3 build profile, SMB/FTP protocol gateways, typed response actions, transactional control, an embedded policy management console, and Helm deployment examples.

— Lists many capabilities but no example invocation shown. v0.16.1

03 Higher default stream reassembly limit IMPROVED

75

The `max-stream-bytes` flag for `net capture` now defaults to 100 MiB (raised from 10 MiB), controlled via `NC_MAX_STREAM_BYTES`.

Cap stream reassembly memory in a 4 GiB container without overriding the cgroup-aware GC headroom Netcap now applies automatically.

```
$ NC_MAX_STREAM_BYTES=52428800 ./net capture -read traffic.pcap -out nc
```

— Exact flag, env var, default change, and runnable example. v0.16.3

04 **TLS 1.2/1.3 decryption via embedded NSS key logs** NEW

70

Netcap now decodes authenticated TLS 1.2 and TLS 1.3 HTTP streams using NSS key logs embedded in pcapng files, and marks decrypted records explicitly.

— Names NSS key log mechanism but no command shown. v0.16.3

05 **Signed database and enrichment management in net util** NEW

65

Install signed enrichment databases with `net util -download-dbs`, refresh them with `net util -update-dbs`, and fetch optional enrichment data with `net util --download-enrichment`.

— Three exact subcommands named for db management. v0.16.1

06 **Gopacket fork adoption and startup reporting** IMPROVED

60

Upstream gopacket is replaced by the `dreadl0ck/gopacket` fork at commit `f2eeabc1ff6e`, adding live-only close wakeups and bounded readers, and Netcap now reports the effective gopacket replacement at startup.

— Names fork and commit but no migration steps. v0.16.3 v0.16.1

07 **Content-based capture file identification** IMPROVED

60

Netcap now identifies capture files by content (magic bytes) rather than `.pcap` / `.pcapng` extension, so inputs with non-standard extensions are accepted by `net capture -read`.

— Explains mechanism and flag but no limits given. v0.16.3

08 **Cgroup-aware memory limits via GOMEMLIMIT** IMPROVED

60

Netcap now respects `GOMEMLIMIT` and applies cgroup-aware Go GC headroom, keeping capture within container memory limits.

— Names env var and behavior, no numeric limits. v0.16.3

THINNER COVERAGE BELOW

09 **New `net split` capture subsetting command** NEW 55

New `net split` command reads PCAP/PCAPNG captures and exports bounded packet subsets and time shards, including time-only exports without a native BPF runtime.

— Describes capability, no flag names or example given. v0.16.1

10 **Live-stream maintenance and flush intervals** NEW 45

Control independent live-stream maintenance and output snapshot intervals with `--live-maintenance-interval` and `--live-flush-interval` on `net capture`.

— Names two flags, no defaults or example given. v0.16.1

11 **Checksum-pinned resumable corpus replay** NEW 45

Netcap adds checksum-pinned resumable corpus replay, allowing interrupted replay runs to continue from a known-good checkpoint.

— Mechanism named, no flag or command given. v0.16.3

12 **SMB DCE/RPC and Netlogon detection coverage** IMPROVED 45

Netcap now recovers additional SMB-carried DCE/RPC operations and covers additional service-control and Netlogon operations in detections.

— Names protocols covered, no flags or examples. v0.16.3

13 **Beacon incident aggregation in c2.beacon** IMPROVED 40

Netcap now aggregates beacon incidents and detects skipped POST polling, extending `c2.beacon` coverage.

— Names c2.beacon area, no thresholds or flags. v0.16.3

14 **New `net mcp` MCP server command** NEW 35

New `net mcp` command exposes a stdio Model Context Protocol server for AI-agent workflows.

— One-line description, no protocol details or example. v0.16.1

15 **NBNS and IEEE 802.15.4 audit records** NEW 30

Netcap adds NBNS and IEEE 802.15.4 audit records, expanding the set of protocol-layer records produced during capture.

— Just names two new record types, no field detail. v0.16.3

16 **SLL2 RARP link-layer frame support** NEW 30

Netcap now handles SLL2 RARP link-layer frames during packet decoding.

— One-line mention, no further detail. v0.16.3

17 **Distinct exit codes for capture outcomes** IMPROVED 30

Netcap now reports distinct exit codes for partial and unsupported exits.

— Brief mention, no actual code values listed. v0.16.3

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ See who can access an AWS asset: NOTES ↗ RANK

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments, helping teams identify and remediate security risks.

> Plerion significantly expanded its AWS inventory coverage, adding collection support for ten service families spanning OpenSearch Serverless, EKS, Route 53 Resolver, Global Accelerator, EventBridge, Aurora DSQL, MSK, Athena, Backup, and IoT Core.

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

01 **Aurora DSQL cluster collection** NEW 73

Plerion now collects Amazon Aurora DSQL clusters, showing status, deletion protection, encryption type and KMS key, endpoint, multi-Region peers, and resource-based policy; requires `dsql` read permissions, present on stacks deployed or updated since July 2026.

— Names exact permission requirement and eligibility window

AWS OpenSearch Serverless in your inven…

02 **OpenSearch Serverless inventory collection** NEW 63

Plerion now collects AWS OpenSearch Serverless collections, VPC endpoints, security configurations, and encryption, network, and data access policies — each showing its type, encryption key, standby replicas, deletion protection, endpoints, tags, and full policy or SAML/IAM federation settings.

— Detailed fields named but passive inventory collection, no user action

AWS OpenSearch Serverless in your inven…

03 **EventBridge connections, destinations, and rule attachment** NEW 61

Plerion now collects Amazon EventBridge connections and API destinations as standalone assets, and attaches up to 300 rules per event bus — each rule showing its event pattern or schedule, state, IAM role, and targets.

— Specific limit (300 rules) named but no direct user action

AWS OpenSearch Serverless in your inven…

04 **MSK VPC connection collection and faster policy updates** IMPROVED 60

Plerion now collects Amazon MSK managed VPC connections, showing the target cluster, VPC, subnets, security groups, authentication type, and state; MSK clusters now also update within minutes when a cluster policy is deleted.

— Names fields and latency improvement, no user action required

AWS OpenSearch Serverless in your inven…

THINNER COVERAGE BELOW

05 **Route 53 Resolver resource collection** NEW 59

Plerion now collects seven Route 53 Resolver resource types: endpoints, resolver rules, query logging configs, DNS Firewall rule groups, domain lists, and the VPC associations of rules and rule groups — including embedded rules, IP addresses, and sharing policies.

— Seven distinct resource types enumerated, no user-triggered action

AWS OpenSearch Serverless in your inven…

06 **Athena workgroup and data catalog collection** NEW 59

Plerion now collects Amazon Athena workgroups and registered data catalogs, showing engine version, query result location and encryption, data scan limit, catalog type, parameters, and status.

— Specific fields like data scan limit named, no action step

AWS OpenSearch Serverless in your inven…

07 **IoT Core resource collection** NEW 55

Plerion now collects six AWS IoT Core resource types: policies, custom authorizers, role aliases, topic rules, CA certificates, and domain configurations — each with its operational settings, status, and tags.

— Six resource types named, no concrete user action

AWS OpenSearch Serverless in your inven…

08 **EKS add-ons and identity resource collection** NEW 53

Plerion now collects four additional Amazon EKS resource types: add-ons, Fargate profiles, identity provider configurations, and Pod Identity associations, each with role, status, OIDC, and tag details.

— Names resource types and fields but no actionable config

AWS OpenSearch Serverless in your inven…

09 **Global Accelerator inventory collection** NEW 53

Plerion now collects AWS Global Accelerator accelerators and their listeners, showing status, static IPs, DNS name, flow-log state, protocol, port ranges, and client affinity.

— Concrete fields listed, passive collection only

AWS OpenSearch Serverless in your inven…

10 **AWS Backup selection collection** NEW 46

Plerion now collects AWS Backup selections, showing the backup plan, IAM role, included and excluded resources, and tag conditions.

— Brief description of fields, no user-facing action

AWS OpenSearch Serverless in your inven…

NEW SUPPORTED PLATFORM AWS asset access analysis via Access tab

> See who can access an AWS asset: NOTES ↗ LATEST STABLE
>
> Plerion adds an Access tab to AWS assets showing inbound/outbound access paths and policy layers (SCPs, RCPs, permission boundaries) as a graph or list.
>
> Open any AWS asset and use the new Access tab to see what can access the asset, what it can access, and who can gain its permissions — viewable as a graph or a list.
>
> Each arrow in the access graph names the permission level it allows (read, write, or administer); hover over an arrow to inspect the policy layers on the path, including identity policies, permission boundaries, SCPs, RCPs, and the resource policy, with the matching statements from the policy document.
>
> The Access tab is rebuilt after each scan.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Google distroless

Feedback ⓘ changes since 2026-11-02

CODE ↗

Distroless is a language focused docker images, minus the operating system.

distroless adds a parallel Go-based sign_and_push runner to replace the generated bash script

└──▷ TRY IT

Run the new parallel sign-and-push pipeline with a custom concurrency level to push, verify, and tag images faster in CI.

```
$ ./bazelisk run :sign_and_push --config=release -- --keyless $KEYLESS
--jobs 32
```

Adds a Go-based `sign_and_push` binary that wraps `crane` and `cosign` and runs image push and offline cosign verification concurrently, controlled by `--jobs` (default 16), against a single initialized `trusted_root.json`.

Attaches image_aspect directly to `sign_and_push` targets and reads digests from OCI `index.json` at runtime, dropping the intermediate `filegroup`, `oci_image_spdx`, and stamp-sensitive `.digest` targets.

└──▷ BREAKING ON UPGRADE

! The intermediate `filegroup`, `oci_image_spdx`, and stamp-sensitive `.digest` Bazel targets are removed; build configs that depended on them directly will break.

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ

SBOM Drift Detection: Catch Changes from Build to Production (2026-11-03)

BLOG ↗ RANK

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

> Anchore Enterprise introduces SBOM drift detection spanning build and runtime: a new `tag_drift` policy gate catches package changes between builds, while the `anchore-k8s-inventory` Helm agent tracks what's actually running in Kubernetes clusters. It also adds SPDX/CycloneDX SBOM export and routes policy evaluation updates to Slack, Jira, GitHub, Teams, email, or webhook.

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

01 Runtime SBOM drift via Kubernetes inventory agent NEW

95

Deploy the `anchore-k8s-inventory` agent via Helm from the `anchore/k8s-inventory` chart at `charts.anchore.io` to poll the Kubernetes API on a configurable interval and report each container's `image_digest`, `image_tag`, `pod_uid`, and cluster name back to Anchore Enterprise. Activating a `runtime_inventory` subscription triggers automatic vulnerability and policy analysis of every image the agent discovers, with findings surfaced in a Kubernetes-aware GUI view broken down by cluster, namespace, node, and pod.

Deploy the runtime inventory agent into a cluster to get continuous visibility into what images are actually running, without requiring broad cluster credentials.

```
$ helm install anchore-k8s-inventory anchore/k8s-inventory --repo
https://charts.anchore.io
```

— Names chart, repo, reported fields, and runnable helm install command

SBOM Drift Detection: Catch Changes fro…

02 tag_drift policy gate for build-time drift NEW

90

Anchore Enterprise adds a `tag_drift` policy gate that detects when a new build changes the packages present in an image compared to a prior build of the same tag. Scope the `packages_modified` trigger to OS-level packages only by setting the `package_type` parameter to `apk`, `rpm`, or `deb`, limiting noise from intentional

language-dependency bumps; a policy rule can set an action such as `STOP` to block the build when drift is detected.

Block a build when any OS package changes between successive builds of the same image tag, using `packages_modified` scoped to `apk` packages.

```json
{
  "action": "STOP",
  "gate": "tag_drift",
  "trigger": "packages_modified",
  "params": [{"name": "package_type", "value": "apk"}],
  "id": "1168b0ac-df6c-4715-8077-2cb3e016cf63"
}
```

— Names gate, trigger, params, and a runnable policy JSON example

SBOM Drift Detection: Catch Changes fro…

THINNER COVERAGE BELOW

03 **Notifications for policy evaluation updates** IMPROVED

45

Send policy evaluation updates — including `tag_drift` trigger results — to Slack, Jira, GitHub, Microsoft Teams, email, or a webhook via Anchore Enterprise Notifications.

— Lists channels but no setup steps or config key

SBOM Drift Detection: Catch Changes fro…

04 **SBOM export in SPDX and CycloneDX** NEW

35

Export SBOMs generated during image analysis in SPDX or CycloneDX formats.

— Names two formats but no export command or endpoint given

SBOM Drift Detection: Catch Changes fro…

WAS THIS USEFUL?  

◆ SaaS Security Posture

Nudge Security

Feedback ⓘ

Detect region-specific sensitive data in AI conversations NOTES ↗

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

Nudge Security adds optional in-browser redaction enforcement so detected sensitive data can be blocked before it leaves the browser.

NEW DETECTION Enable one-click in-browser redaction of detected sensitive data, with optional enforcement so unredacted data never leaves the browser, via Settings > In-browser nudges.

Choose which regional data types to monitor by selecting active regions in Browser Extension Settings.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ Secure SaaS - November 9, 2026 BLOG ↗

RANK

Sysdig Secure is a cloud-native security platform that detects threats, monitors runtime activity, and scans container images and Kubernetes environments for vulnerabilities and misconfigurations.

> Sysdig Secure's Registry Scanner now supports credential-free authentication to Azure Container Registry on AKS, along with clearer diagnostics for identity-related scan failures.

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

01 Managed/workload identity auth for ACR on AKS NEW

70

Registry Scanner now authenticates to Azure Container Registry using managed identity or workload identity on AKS, via the cluster's kubelet identity or a user-assigned identity federated to its service account, removing the need to store registry credentials. Both RBAC and RBAC + ABAC repository permission modes are supported.

— Explains mechanism and scope but no config steps given Secure SaaS - November 9, 2026

THINNER COVERAGE BELOW

02 ACR identity authentication failure logging IMPROVED

45

Registry Scanner adds specific log messages for ACR identity authentication failures, distinguishing between an identity that cannot get a token, one with no role on the registry, and requests blocked by registry network rules.

— Names failure categories but no log format or location Secure SaaS - November 9, 2026

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SOAR

Torq

Feedback ⓘ changelog · Nov 10, 2026 CHANGELOG ↗ RANK

Transform the enterprise SOC with Torq — combining agentic SecOps and Hyperautomation to triage, investigate, and respond to threats faster than ever.

> Torq enhanced case management with Markdown-formatted case tabs and brought one-click quick actions directly to observables across multiple surfaces.

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

01 **Quick actions on observables** NEW 60

Quick actions now appear directly on observables — attach a workflow to any observable and it surfaces on the observable's own page, the case's Observables tab, and the hover card.

— Names three exact surfaces where action appears changelog-20261110-ed912edc

THINNER COVERAGE BELOW

02 **Markdown sections in custom case tabs** NEW 45

Custom case tabs now support free-form Markdown sections alongside custom fields, letting workflows populate each case tab with tables, formatted text, and structured reports.

— Describes capability but no config key or command shown changelog-20261110-ed912edc

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

pysigma

Feedback ⓘ v2.1.0 NOTES ↗ RANK

Pysigma is a python library to parse and convert Sigma rules into queries (and whatever else you could imagine)

> pySigma adds a new module for matching events against Sigma rules directly in Python, without needing to convert rules to a query language first.

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

01 **Event matching via sigma.matching module** NEW 55

The new `sigma.matching` module lets users match live or logged events directly against Sigma rules in Python code, without first converting those rules to a query language.

— Names the module but no API details or usage example v2.1.0

NEW SUPPORTED PLATFORM ADFS Windows logsource

v2.0.3 NOTES ↗ STABLE

pySigma adds optional strict mode for field mapping transformations, preserves finalizer configs on pipeline load, and allows dashes in `1/all of` selector patterns.

└──▷ GET THIS VERSION ▾

pip

```
$ pip install pysigma==2.0.3
# already installed? upgrade to this version:
$ pip install --upgrade pysigma==2.0.3
```

git

```
$ git clone --branch v2.0.3 https://github.com/SigmaHQ/pySigma.git
# already have the repo? check out this version:
$ git checkout v2.0.3
```

PyPI attestation · checked 10 Nov Wrong command?

Add an optional strict mode for field mapping transformations to catch unmapped fields at conversion time rather than silently passing them through.

Preserve finalizer configurations when loading processing pipelines, so pipeline-defined finalizers are no longer silently dropped on load.

> Allow dashes in `1/all of` selector patterns so rule conditions referencing hyphenated detection identifiers now parse correctly.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261109.0 NOTES ↗ RANK

OpenCTI is an open Cyber Threat Intelligence Platform

> OpenCTI's biggest addition this window is a platform-wide custom fields system that lets administrators extend entities and relationships with typed attributes usable across filters, CSV/JSON mappers, widgets, exports, and full-text search; alongside that it shipped Form Intake submission capabilities, configurable per-relationship deduplication windows, TLS/GraphQL security hardening, unified Filigran XTM Platform licensing, and smaller improvements to user management, dashboards, and date/time display.

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

01 **Custom fields for entities and relationships** NEW 96

Administrators manage custom fields (create, edit, delete, attach to entity types) in `Settings > Taxonomies > Custom fields` using the Manage custom fields capability, set per-entity-type defaults using `@now` for date fields, and mark fields mandatory (bypassed by Bypass mandatory attributes). Technical names start with `x_opencti_cf_` (e.g. `x_opencti_cf_priority`) and are usable in the API, filters, exports, and STIX bundles; fields can be mapped at ingestion via aliases (STIX custom properties starting with `x_`), imported from STIX bundles, CSV mappers, and JSON mappers by technical name or alias, and exported in STIX exports/streams as properties in the object's OpenCTI extension. Custom fields work as filter keys in lists, dashboards, saved filters, streams, triggers, and playbooks with the same operators as standard attributes, as columns in list widgets when filtered to a single entity type, and in attribute widgets in Fintel templates and custom views; values can be bulk-updated via the Manipulate knowledge playbook component or the Update background task action. Seven value types are supported — Text, Markdown, Integer (with min/max bounds), Boolean, Date, Selection list, and Multiple selection list — attachable to Stix core relationship (covering every relationship type) or to Sighting; custom field values are also searchable via the platform's full-text search and a custom resolution field can be added to cases through this system. Without an Enterprise Edition license, existing values remain readable, filterable, importable, and exportable, but creating or managing fields requires the license.

Ingest a STIX bundle from a connector that uses a proprietary priority property, mapping it to a custom field via an alias so no connector-side changes are needed.

> 📍 Define the custom field with technical name `x_opencti_cf_priority` and add `x_mycompany_priority` as an alias in `Settings › Taxonomies › Custom fields`. The bundle property `x_mycompany_priority` will then populate the custom field on ingest automatically.

Set a creation-date default on a custom date field so every new entity gets a timestamp without manual entry.

> 📍 In `Settings › Customization › Entity types`, select the target entity type, go to its `Attributes` tab, attach the date custom field, then set its default value to `@now` using the edit icon.

— Names config path, capability, prefix, value types and config example `7.261109.0`

---

02 **TLS certificate verification controls** NEW 93

New TLS verification settings govern outgoing connections platform-wide: `tls:reject_unauthorized` / `TLS__REJECT_UNAUTHORIZED` is a default for feeds, synchronizers, SMTP, RabbitMQ, HTTPS proxy, and Filigran services (opt in early by setting it `true` before it becomes the default); `https_proxy_reject_unauthorized` / `HTTPS_PROXY_REJECT_UNAUTHORIZED` independently overrides the HTTPS proxy and otherwise inherits from `tls:reject_unauthorized`; and `opencti:tls_reject_unauthorized` / `OPENCTI_TLS_REJECT_UNAUTHORIZED` sets the default for workers and connectors, inherited by `opencti:ssl_verify` / `OPENCTI_SSL_VERIFY` and by `worker:listen_api_ssl_verify` for connectors in HTTP mode.

Opt in to strict TLS verification for all platform outgoing connections before it becomes the default, while also hardening GraphQL in one config block.

```yaml
tls:
  reject_unauthorized: true

app:
  graphql:
    secure_by_default: true
    batching_protection:
      query_default: 10
      mutation_default: 5
```

— Every config key and inheritance path named with a config example `product docs`

---

03 **Relationship deduplication window configuration** `BREAKING`

88

Configure the deduplication merge window globally or per relationship type using the `relations_deduplication` config key, with `past_days` and `next_days` accepting a number of days or an ISO 8601 duration (e.g. `PT30M`, `PT6H`, `P1DT12H`), or override per type via environment variables such as `RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__PAST_DAYS` and `RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__NEXT_DAYS` set on every platform node. When an incoming entity carries no confidence level, the existing entity keeps its own confidence, or receives the sending user's maximum confidence if it has none. A per-type override replaces the global configuration entirely — any omitted `past_days` / `next_days` key resolves to 0, not the global default — and invalid duration values (e.g. `30m` instead of `PT30M`) are silently ignored, also producing a window of 0 with no validation error.

Narrow the sighting deduplication window to 30 minutes to avoid merging sightings that are close in time but distinct events.

```json
"relations_deduplication": {
  "types_overrides": {
    "stix-sighting-relationship": {
      "past_days": "PT30M",
      "next_days": "PT30M"
    }
  }
}
```

Apply the same 30-minute sighting window via environment variables across all platform nodes without touching the config file.

```
$ export RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__PAST_DAYS=PT30M
export RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__NEXT_DAYS=PT30M
```

Narrow the sighting deduplication window to ±30 minutes to avoid merging sightings of distinct events close in time.

```json
"relations_deduplication": {
  "types_overrides": {
    "stix-sighting-relationship": {
      "past_days": "PT30M",
      "next_days": "PT30M"
    }
  }
}
```

Apply the same ±30-minute sighting window without touching config files, by exporting environment variables on every platform node.

```
$ export RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__PAST_DAYS=PT30M
export RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__STIX-SIGHTING-RELATIONSHIP__NEXT_DAYS=PT30M
```

— Runnable config/env examples, but silent-failure behavior is the real risk product docs

04 **GraphQL security hardening defaults** NEW 88

`app:graphql:secure_by_default` / `APP__GRAPHQL__SECURE_BY_DEFAULT` enables a secure posture that turns on GraphQL Armor, strips suggestions from error messages, disables the playground and introspection, and restricts WebSocket to subscriptions only. Alias-based batching protection caps are configurable via `app:graphql:batching_protection:query_default` / `APP__GRAPHQL__BATCHING_PROTECTION__QUERY_DEFAULT`, `app:graphql:batching_protection:query_subtypes` / `APP__GRAPHQL__BATCHING_PROTECTION__QUERY_SUBTYPES`, and `app:graphql:batching_protection:mutation_default` /

`APP__GRAPHQL__BATCHING_PROTECTION__MUTATION_DEFAULT`, with the `token` mutation always hard-limited to 1.

Opt in to strict TLS verification for all platform outgoing connections before it becomes the default, while also hardening GraphQL in one config block.

```yaml
tls:
  reject_unauthorized: true

app:
  graphql:
    secure_by_default: true
    batching_protection:
      query_default: 10
      mutation_default: 5
```

— Names every config key and cap; example config block provided `product docs`

05 **Configurable date and time display formats** NEW 86

Platform parameters now include Date format in tables (Default or ISO 8601 `YYYY-MM-DD`) and Time in tables (Default, Date only, `HH:mm`, or `HH:mm:ss`), controlling how dates appear across every table and list platform-wide. Individual users can override this in their profile under User experience using the same two settings, with a Platform default option that defers to the platform setting.

Set the platform to display unambiguous ISO 8601 timestamps with seconds in every table — useful for a SOC sharing incident lists with external customers.

📍 In the platform administration, go to Parameters › (platform settings card) › set Date format in tables to ISO 8601 (YYYY-MM-DD) and Time in tables to Hours, minutes and seconds (HH:mm:ss). Dates will appear as e.g. 2026-10-27 15:32:17+01:00 across all entity and relationship lists, audit logs, alerts, and connector works.

Override the platform date format for your own account without affecting other users.

📍 Click the top-right profile icon › Profile › User experience card › set Date format in tables and Time in tables to your preferred values. Choosing any value other than Platform default takes precedence over the platform setting.

— Exact UI navigation and setting values with sample output format `product docs`

06 **Form Intake submission capabilities** NEW 77

A `Create knowledge via Form Intake` capability lets users submit Form Intakes to create knowledge without the general import-knowledge capability; submissions always create new records rather than updating existing ones, and only users who also hold `Create / Update knowledge` in the destination can update matched existing objects. A draft-scoped variant, 'Create knowledge via Form Intake (creation forced to draft)', forces all submissions into draft mode for restricted users. While filling in a form, users can create entities, observables, labels, and external references on the fly (with attached files and markdown images), held in the browser and created together with the rest of the submission only on submit; submissions are processed asynchronously, with a success message confirming the submission has been queued rather than that objects already exist.

Grant an analyst the ability to submit intelligence forms in draft only, so their intake goes into a review queue rather than directly into the main knowledge base.

> 📍 In OpenCTI, go to Settings › Security › Users (or the relevant Role), then enable the capability `Create knowledge via Form Intake` scoped to draft mode. The user can now submit Form Intakes, and all submissions will land in a draft for review before promotion.

– Capability names and before/after given; UI path only sketched `product docs`

---

07 **Elasticsearch 9 client selection** BREAKING 71

Select the Elasticsearch/OpenSearch client explicitly with `elasticsearch:engine_client` / `ELASTICSEARCH__ENGINE_CLIENT`, accepting `elasticsearch8`, `elasticsearch9`, or auto-detection; this replaces the now-ignored `engine_check` configuration key, so deployments relying on `engine_check` must migrate.

Pin OpenCTI to the Elasticsearch 9 client when a proxy prevents the platform from reading the server version at startup.

```yaml
elasticsearch:
  engine_client: elasticsearch9
```

– Exact replacement key and migration example given, impact scope thin `product docs`

---

08 **Unified Filigran XTM Platform licensing** NEW 69

Paste a Filigran XTM Platform license PEM into the existing license field (settings page or `APP__ENTERPRISE_EDITION_LICENSE`) to activate Enterprise Edition across

OpenCTI, OpenAEV, and OpenCRQ modules. Instances can also register with XTM One to activate Enterprise Edition without a local license; OpenCTI verifies the Filigran-signed license at every registration and shows Filigran license (through XTM One) as the license source on the settings page.

Activate Enterprise Edition on a self-hosted OpenCTI instance by supplying the Filigran XTM Platform license via environment variable instead of the UI.

```env
APP__ENTERPRISE_EDITION_LICENSE=<filigran-pem-certificate>
```

Activate Enterprise Edition on an air-gapped OpenCTI instance by supplying a Filigran XTM Platform license via environment variable.

```env
APP__ENTERPRISE_EDITION_LICENSE=<filigran-pem-certificate>
```

— Env var and registration flow given, but activation internals unexplained `product docs`

THINNER COVERAGE BELOW

09 **Prometheus metrics endpoint host binding** NEW — 48

Bind the Prometheus metrics endpoint to a specific IPv4 address using `app:telemetry:metrics:exporter_prometheus_host` / `APP__TELEMETRY__METRICS__EXPORTER_PROMETHEUS_HOST`.

— Exact config key given but no example or use case `product docs`

10 **Merge duplicate user accounts and lockout policy** NEW — 42

Administrators can merge duplicate user accounts — for example, merging a local account into an SSO-created account — via the Merging users workflow in the user management interface, and can configure account lockout policy settings to control how repeated failed logins are handled.

— Names the workflow but gives no exact settings path or thresholds `7.261109.0`

11 **Time-series breakdown dimension in dashboard widgets** NEW — 42

The dashboard widget editor now supports configuring a time-series breakdown dimension, letting a time series be split by a chosen attribute in a single chart; this breakdown is supported in both private and public dashboards.

— Describes capability but no config path or example attribute `7.261109.0`

12 **Enrichment connector launch from investigation graph** NEW

39

Run an enrichment connector directly from the investigation graph's right-click menu without leaving the graph view.

— UI entry point named but no connector list or mechanism 7.261109.0

13 **Ingestion health monitoring enhancements** IMPROVED 34

Ingestion health monitoring now tracks last-run failures, feed and sync access-denied events, errors from ping, and disabled users.

— Lists tracked signals only, no UI path or alerting mechanism 7.261109.0

14 **Label export/import in configuration presets** NEW 26

Labels can now be exported and imported as part of configuration presets.

— One-line statement with no format, path, or scope detail 7.261109.0

└──▷ ALSO FROM THESE RELEASES

Override the deduplication window for sighting relationships to 30 minutes while keeping the platform default for all other types.

```env
RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__stix-sighting-
relationship__PAST=PT30M
RELATIONS_DEDUPLICATION__TYPES_OVERRIDES__stix-sighting-
relationship__NEXT=PT30M
```

└──▷ BREAKING ON UPGRADE

! A relationship type entry in `relations_deduplication.types_overrides` replaces the global configuration entirely: any `past_days` or `next_days` key omitted from the override resolves to 0, not the global default.

! Invalid `past_days` or `next_days` values (e.g. `30m` instead of `PT30M`) are silently ignored and produce a window of 0, with no validation error.

! A relationship type override in `relations_deduplication.types_overrides` replaces the global configuration entirely: any key (`past_days` or `next_days`) absent from the override resolves to 0, not the global default.

! Invalid duration values (e.g. `30m`, which is neither a plain number nor a valid ISO 8601 duration) are silently ignored and produce a window of 0 — no validation error is raised.

! The `engine_check` Elasticsearch configuration key is now ignored; replace it with `elasticsearch:engine_client` / `ELASTICSEARCH__ENGINE_CLIENT`.

WAS THIS USEFUL?  

◆ SIEM

Wazuh

Feedback ⓘ v4.14.10 NOTES ↗ RANK

Wazuh - The Open Source Security Platform. Unified XDR and SIEM protection for endpoints and cloud workloads.

> Wazuh's latest updates focus on operational safeguards: memory thresholds that cap runaway rulesets in analysisd, queue-size warnings for the indexer connector and vulnerability scanner, new RBAC read-only roles, event logging in plaintext and JSON, and Puppet module support for managing indexer credentials.

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

01 **Rule tree memory thresholds in analysisd** NEW 95

Set `analysisd.rule_tree_memory_warning` in the internal options file (or `<rule_tree_memory_warning>` in `ossec.conf`) to log a warning when the rule tree exceeds a memory threshold during ruleset loading; accepts a byte size (e.g. `2G`) or a percentage of physical or cgroup memory (e.g. `25%`), defaulting to `25%`, with `0` disabling it. Set `analysisd.rule_tree_memory_limit` (same size/percentage format, disabled by `0` default) to reject a ruleset that would exceed a hard memory cap, blocking uploads via the Wazuh server API and dashboard with a descriptive error.

Cap rule tree memory on a container-based deployment so that an oversized ruleset upload is rejected before it exhausts cgroup memory.

ini

```
analysisd.rule_tree_memory_warning=20%
analysisd.rule_tree_memory_limit=50%
```

Prevent an oversized ruleset from consuming all available memory by capping rule-tree growth and getting an early warning before the hard limit is hit.

xml

```
<ossec_config>
  <analysisd>
    <rule_tree_memory_warning>75%</rule_tree_memory_warning>
    <rule_tree_memory_limit>90%</rule_tree_memory_limit>
  </analysisd>
</ossec_config>
```

— Names exact config keys, formats, defaults and a runnable config example. v4.14.10

02 **Puppet module support for indexer credentials** BREAKING

80

The `wazuh::manager` Puppet class gains `vulnerability_indexer_password` and `vulnerability_indexer_username` parameters to configure the Wazuh manager keystore on every Puppet run and to override the default indexer account. Because the password is written to the keystore on every run, it must match the current indexer password or the manager becomes misconfigured after each run.

Set the indexer password and a custom indexer username in your Puppet manifest so the Wazuh manager keystore stays in sync on every run.

puppet

```
class { 'wazuh::manager':
  vulnerability_indexer_password => 'MySecretPass123',
  vulnerability_indexer_username => 'custom-indexer-user',
}
```

— Names exact parameters, class, and a runnable manifest, plus the breaking caveat.

product docs

03 **Read-only RBAC roles for Reporting, Notifications, Snapshot Management**

NEW

65

Map users to the new `reports_read_access`, `notifications_read_access`, or `snapshot_management_read_access` roles in the Wazuh RBAC interface to grant read-only access to Reporting, Notifications, and Snapshot Management without exposing dashboard write permissions (avoid mapping to `kibana_user` or `kibanauser` backend roles).

Grant a read-only user access to Reporting, Notifications, and Snapshot Management without exposing dashboard write permissions.

> 📍 Map the user to the following roles in the Wazuh RBAC interface: - reports_read_access - notifications_read_access - snapshot_management_read_access Do NOT map to: kibana_user role or kibanauser backend role.

— Names exact roles and UI steps but no deeper mechanism. product docs

THINNER COVERAGE BELOW

04 **TCP session counter in wazuh-remoted** NEW

50

`wazuh-remoted` statistics now include a `tcp_sessions_unassociated` counter reporting TCP connections not yet associated with an agent.

— Names the exact stat field but no usage guidance. v4.14.10

---

05 **Queue-size warnings for indexer connector and vulnerability scanner**

NEW 50

Wazuh now logs a warning when a queue of the indexer connector or of the vulnerability scanner holds 100,000 elements or more, and again each time its size doubles.

— Gives concrete threshold and repeat behaviour but no config surface. v4.14.10

---

06 **Plaintext and JSON event logging formats** NEW 40

Store logs received from third-party applications and systems via event logging in either plaintext or JSON format; both are disabled by default.

— States what the formats do but lacks config detail or commands. product docs

└──▷ **BREAKING ON UPGRADE**

! The `wazuh::manager` class now writes `vulnerability_indexer_password` to the Wazuh manager keystore on every Puppet run; the value must match the current indexer password or the manager will be misconfigured after each run.

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Databricks Genie app for Slack is now generally available (2026-11-09) BLOG ↗

RANK

Databricks is a cloud data lakehouse platform that unifies data engineering, analytics, machine learning, and AI workloads on Apache Spark.

> Databricks shipped a wide range of updates this window: AI Function tasks now run natively inside Lakeflow Jobs, Designer gained a mechanism-rich Multi-Row Formula operator and file export to Google Drive/SharePoint, Genie expanded with Ontology (GA), Clean Rooms, Volumes support and Slack/Teams approvals, and Unity Catalog added volume recovery (`UNDROP VOLUME`) and native Secrets management, alongside Asset Bundles' new `sync_mode` deployment controls and Unity Gateway's desktop AI agent provisioning.

└──▷ **WHAT SHIPPED · 29 FEATURES** most completely described first ? what's the number?

01 **Multi-Row Formula operator in Designer** NEW 100

Designer adds a Multi-Row Formula operator that computes new columns from neighboring rows using shared Sort and Grouping fields to control row windows and group boundaries. It supports Running aggregate (Sum, Average, Minimum, Maximum, Count), `Row number / rank` (sequential, gap-rank, dense-rank, with Start at), Value from another row (lag/lead via Rows away and If no such row), Difference from previous row (via Rows back), Moving average over a configurable Window, Exponential smoothing with a Smoothing factor (α, default 0.3), and Custom expression for arbitrary Spark SQL `OVER` window expressions.

Compute a per-customer running total and row-over-row change in a visual workflow without writing SQL.

📍 In the Designer canvas, add a Multi-Row Formula operator after your source, then: 1. Click Add formula › Running aggregate . Set New column name to `customer_running_total` , Function to Sum , Of column to amount . Add `order_date` (Ascending) as Sort and `customer_id` as Grouping. 2. Click Add formula › Difference from previous row . Set New column name to `amount_change` , Of column to amount , Rows back to 1. Same Sort and Grouping. 3. Click Apply .

— Every formula mode named with its controls and a runnable walkthrough product docs

02 **`workspace.sync_mode` deployment control in Asset Bundles**

NEW 95

Control how bundle files reach deployed resources with the new `workspace.sync_mode` config key (from version 1.23.0), accepting `overwrite`, `in-place`, or `immutable` — defaulting to `in-place` for development deployments from the workspace and `overwrite` otherwise. Setting `workspace.sync_mode: immutable` uploads bundle files and artifacts as a single immutable snapshot to block non-admin edits, recoverable with `databricks bundle deploy --force` after an out-of-band change; `workspace.sync_mode: in-place` references workspace source files directly without uploading. This replaces the deprecated `presets.source_linked_deployment: true` setting (use `in-place` instead of `true`, `overwrite` instead of `false`).

Prevent non-admin users from modifying deployed production job files by using an immutable snapshot deployment.

```yaml
sync_mode: immutable
```

Force a fresh immutable snapshot after a snapshot was modified outside the bundle, unblocking the next deployment.

```
$ databricks bundle deploy --force
```

— Names every mode, default, recovery command and migration path `product docs`

---

03 **UNDROP VOLUME and dropped-volume recovery** NEW 95

Restore a dropped volume with `UNDROP VOLUME` and list dropped volumes with `SHOW VOLUMES DROPPED`. On Databricks Runtime 19 or above, recover a dropped Unity Catalog volume within a 7-day window by locating its ID via `SHOW VOLUMES DROPPED` and restoring it — along with its original privileges and tags — using `UNDROP VOLUME WITH ID`.

Recover a volume accidentally dropped, then confirm it is no longer listed as dropped.

```
$ SHOW VOLUMES DROPPED;
UNDROP VOLUME <volume_name>;
```

Restore an accidentally dropped Unity Catalog volume, preserving its original privileges and tags, within the 7-day recovery window.

```
$ SHOW VOLUMES DROPPED;
-- Note the volume_id from the output, e.g. 'abc123'
UNDROP VOLUME WITH ID 'abc123';
```

— Exact SQL commands, recovery window and runtime requirement given

Recover dropped Unity Catalog volumes w…

04 **`UCDataset` for loading Unity Catalog tables into PyTorch**

NEW 85

Load Unity Catalog tables directly into PyTorch training loops on AI Runtime using `UCDataset`, with no Spark or SQL warehouse required. Apply filter expressions and deletion vectors at read time so only matching rows are streamed, and `UCDataset` automatically splits matching rows across GPUs and data loader workers (Beta).

— Names the class and its read-time and distribution behavior

Load Unity Catalog tables into PyTorch …

05 **Desktop AI agent provisioning via Unity Gateway** NEW

80

Configure Claude desktop and ChatGPT desktop work agents with governed models, MCP servers, and skills through Unity Gateway, with subscription relay available in Claude desktop. Provision the Claude desktop app machine-wide via MDM software using `ug configure --global` (macOS only), letting users import their claude.ai projects and chats.

Deploy a governed Claude desktop configuration to all managed Macs in your org without requiring users to configure it themselves.

```
$ ug configure --global
```

— Exact MDM command given alongside what it configures

Provision Claude desktop with MDM using…

06 **Secrets in Unity Catalog** NEW

80

Databricks now supports Secrets in Unity Catalog, letting you create secrets, grant access, and read them in notebooks, jobs, and UDFs. Grant broad secrets discoverability in a catalog — without exposing secret values or allowing code references — using `GRANT BROWSE ON CATALOG <catalog_name> TO <principal>;`.

Grant broad secrets discoverability in a catalog without exposing secret values or allowing code references.

```
$ GRANT BROWSE ON CATALOG <catalog_name> TO <principal>;
```

— Names usage surfaces plus a runnable GRANT statement product docs

07 **User-aware workspace instructions in Genie** NEW 75

Customize Genie chat behavior per user group by adding group-specific logic to `.genie_workspace_instructions.md`, including a `Load the user's info` instruction before group-specific rules.

Route Finance and Sales users to different metrics in Genie chat responses without changing their data permissions.

```markdown
# User-aware behavior
- Load the user's info to get their group memberships.
- If the user belongs to the Finance group, include gross margin and operating expenses in financial summaries.
- If the user belongs to the Sales group, include bookings and pipeline in performance summaries.
```

— Names the exact config file and shows a working instruction block product docs

08 **AI Functions runtime requirement raised to DBR 17.3 LTS** 75

BREAKING

Starting January 2027, AI Functions that previously required DBR 15.4 LTS or above (except `ai_similarity`) will require Databricks Runtime 17.3 LTS or above when called from SQL queries; calls on runtimes below 17.3 return an error, while `ai_similarity` and the AI Functions REST API are unaffected.

— Clear migration date and version threshold, no command needed product docs

09 **Okta connector in Lakeflow Connect** NEW 75

A new Okta connector in Lakeflow Connect (Beta) ingests identity and access data — users, groups, group memberships, app integrations, devices, policies, and administrator roles — using SSWS API token authentication with incremental ingestion.

— Lists all ingested object types and the auth method product docs

10 **History control on shared views via ALTER SHARE** NEW 75

Control history sharing per view by adding `WITH HISTORY` or `WITHOUT HISTORY` to `ALTER SHARE ADD VIEW` and `ALTER SHARE ALTER VIEW` statements; existing

history defaults are unchanged.

Explicitly exclude history when sharing a view in a Delta Sharing share, so recipients see only the current snapshot.

```
$ ALTER SHARE my_share ADD VIEW my_catalog.my_schema.my_view
WITHOUT HISTORY;
```

— Exact SQL syntax and a runnable example given `product docs`

---

11 **Date arithmetic in Lakeflow Jobs parameters** NEW 75

Shift a date in Lakeflow Jobs parameters by a duration before formatting it using dynamic value references such as `{{(job.start_time - 1.day).iso_date}}` (Beta).

Pass yesterday's date to a task parameter — useful for incremental ETL jobs that should process the previous day's data.

```text
{{(job.start_time - 1.day).iso_date}}
```

— Exact reference syntax with a practical ETL example

`Shift dates in Lakeflow Jobs parameters…`

---

12 **AI Function task type in Lakeflow Jobs** NEW 70

Lakeflow Jobs gains an AI Function task type (Beta) that runs `ai_parse_document`, `ai_extract`, or `ai_classify` on volume files or table rows without writing SQL; each run processes only inputs that are new or changed since the previous run.

— Names functions and incremental behavior, no exact config shown `product docs`

---

13 **Workflow export to Google Drive and SharePoint** NEW INTEGRATION

NEW 70

Write Designer workflow output files to Google Drive or SharePoint by setting `Destination` to `Google Drive` or `Microsoft 365`, entering a `Folder URL`, and signing in via the connector.

Export workflow results to a SharePoint folder for sharing with stakeholders who don't have Databricks access.

> 📍 In the Designer canvas, connect your final operator's outgoing port to a Destination operator. Set `Destination` to `Microsoft 365`, enable the SharePoint connector and click `Connect`. Paste your SharePoint folder link into `Folder URL`, enter a `File name`, and choose Excel as the format. Click `Apply`.

– Exact fields and a step-by-step example make it directly usable product docs

14 **MCP connectors enabled by default for compliance workspaces**

NEW INTEGRATION BREAKING

65

Databricks-provided MCP connectors in Agent (Google Drive , Gmail , Google Calendar , Microsoft 365 , Atlassian , Slack , and GitHub) will be enabled by default starting December 2026 for workspaces with the compliance security profile enabled.

– Names all affected connectors and the exact trigger date product docs

THINNER COVERAGE BELOW

15 **GitHub Action skill publishing to Unity Gateway** NEW

55

Publish Git-managed skills to Unity Gateway using a GitHub Action (beta): map folders to skills and trigger updates on merges to your publishing branch.

– Describes the mechanism but no exact action name or config product docs

16 **Moonshot AI Kimi K3.1 Pro on Unity Gateway** NEW

55

Unity Gateway now supports Moonshot AI Kimi K3.1 Pro (with vision, reasoning, and agentic tool use) on pay-per-token endpoints in `us-east-1` .

– Names model, capabilities and region but no usage steps product docs

17 **AI Functions spend tracked in Unity Gateway budgets** IMPROVED

55

Unity Gateway budgets now include spend from AI Functions such as `ai_parse_document` , `ai_extract` , and `ai_classify` across batch and real-time execution paths.

– Names the tracked functions but no budget config shown product docs

18 **RocksDB default state store in Runtime 17.3 LTS** IMPROVED

55

Databricks Runtime 17.3 LTS and above now defaults Structured Streaming to the RocksDB state store provider, with changelog checkpointing also enabled by default.

– Names the default change but no opt-out or tuning detail product docs

19 **Genie Ontology (GA)** NEW

55

Genie Ontology is now generally available: a business-aware context layer for Genie One, Genie Code, and other AI agents, combining Unity Catalog semantics you define with knowledge Genie extracts from your assets and usage patterns.

— Explains purpose and GA status but no configuration steps

Genie Ontology is now generally availab…

20 **Ingress allowlist for BI tools and Genie chat apps** NEW

50

Allow traffic from Power BI, Tableau Cloud, dbt Cloud, and the Genie apps for Slack and Microsoft Teams in context-based ingress policies.

— Names all allowed tools but not the policy syntax product docs

21 **Power BI semantic models as Genie knowledge source** NEW

50

Pull Power BI semantic model metadata into Genie as external knowledge sources (Beta), enabled by an account admin from the account console.

— Names the admin entry point but no detailed steps product docs

22 **Read replicas for OLTP Data API reads** NEW 45

Databricks OLTP now supports read replicas, letting you offload Data API read requests from primary compute by sending them to a read replica's URL.

— Names the capability but gives no setup steps or limits product docs

23 **Genie Clean Room (Beta)** NEW 45

A new Genie Clean Room (Beta) collaboration type lets a provider share a curated Genie Agent with one collaborator, who can ask it questions in natural language without accessing the underlying data, context, or proprietary IP.

— Describes the model but no setup or API detail product docs

24 **Cross-workspace upstream job triggers in Lakeflow** NEW

45

Trigger a Lakeflow Job when one or more upstream jobs complete, including across workspaces attached to the same metastore (Beta).

— Describes the trigger condition but no setup path product docs

25 **Expandable visualization side panel in Genie** NEW 40

Open any visualization in a side panel at larger size by clicking the expand control in the lower-right corner of the visualization card.

— Gives exact UI location but no further mechanism product docs

26 **Genie tool action approvals from Slack and Teams** NEW

40

Approve or deny Genie tool actions directly from Slack and Microsoft Teams approval prompts without leaving the chat (Public Preview).

— States the capability with no configuration detail product docs

27 **Named feature flags in Asset Bundles** NEW 40

Define named feature flags for a bundle using the new `features` mapping, where each key is the name of a feature (available from version 1.22.0).

— Names the config key but no usage example product docs

28 **Genie support for Volumes** NEW 40

Databricks Genie now supports Volumes so users can ask questions about unstructured files alongside structured data (Public Preview).

— States the capability with no mechanism detail product docs

29 **Tableau and Sigma as planned Genie knowledge sources** NEW

20

Databricks plans to support Tableau and Sigma as external knowledge sources for Genie.

— Roadmap mention only, no mechanism or timeline product docs

└──▷ **ALSO FROM THESE RELEASES**

.
.
.
.

└──▷ **BREAKING ON UPGRADE**

! Beginning January 2027, AI Functions that previously required DBR 15.4 LTS or above (except `ai_similarity`) now require DBR 17.3 LTS or above; calls on runtimes below 17.3 return an error.

! AI Functions in SQL queries will no longer be supported on Databricks Runtime versions below 17.3 LTS starting January 2027; `ai_similarity` and the AI Functions REST API are unaffected.

! `presets.source_linked_deployment` is deprecated; set `workspace.sync_mode: in-place` instead of `source_linked_deployment: true`, or `workspace.sync_mode: overwrite` instead of `source_linked_deployment: false`.

! AI Functions will require Databricks Runtime 17.3 LTS or above (an upcoming change from the current 15.4 LTS minimum).

> Databricks Genie app for Slack is now generally available (2026-11-09) BLOG ↗
>
> STABLE

Genie Search is now generally available (2026-11-09) BLOG ↗ STABLE

Databricks adds Genie Search (GA), indexing org knowledge across connected sources with permission-aware results inside Genie One.

Databricks adds Genie Search, a generally available knowledge-indexing capability that connects multiple sources so Genie One can search all indexed content in one place.

Genie Search indexes each document's permissions alongside its content, so users see only documents they can already access in the source system.

Genie One now includes an org-directory feature powered by Genie Search that lets users find teammates, browse the org chart, and view profiles sourced from Workday.

Genie One's Discover view now supports searching Databricks assets and content from connected sources in a single search experience via Genie Search.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

Connection Policies can now manage how agents connect to Smart Protection Services (2026

BLOG ↗ RANK

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

> Trend Vision One shipped new Connection Policy controls for routing Smart Protection Services traffic, custom rule authoring across Integrity Monitoring, Intrusion Prevention, and Log Inspection policy resources, and richer attack-mode and impact visibility for AI Scanner findings.

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

01 Connection Policies route Smart Protection Services traffic

NEW 95

Runtime Connection Policy controls under Endpoint Security → Endpoint Security Configuration → Connection Policies let administrators route Smart Protection Services traffic through selected Service Gateways, any available Service Gateway, or directly to Trend Vision One. This requires agent version 202610 or later and a Service Gateway running Smart Protection Services 2.0.11 or later; enrollment applies to ALL agents assigned a connection policy and cannot be undone, so policies must be configured before enrolling.

Configure agent traffic routing before enrolling in Smart Protection Services to avoid an irreversible policy state.

📍 In the Trend Vision One console, go to Endpoint Security → Endpoint Security Configuration → Connection Policies, open or create a runtime connection policy, set the Smart Protection Services connection method to selected Service Gateways, 'any available Service Gateway', or 'directly to Trend Vision One', then enroll in Smart Protection Services.

— Names exact UI path, version requirements, and irreversible enrollment behavior.

Connection Policies can now manage how …

02 Attack mode and impact columns in AI Scanner findings

IMPROVED 78

The AI Scanner's Evaluation results table gains a Mode column showing which attack mode (Managed , Custom , or Agentic) produced each finding, also displayed beside the attack objective in Scan details, plus an Impact column with plain-language labels

such as `Data disclosed: High`, `Data tampered: High`, or `Service disrupted: Low`. The Objective filter was updated to group objectives under each attack mode, so selecting a mode and an objective returns only that mode's findings.

Narrow a review session to only findings produced by the Agentic attack mode so you can triage that mode's results in isolation.

> 📍 In Trend Vision One, go to AI Security → AI Application Security → Evaluation results. Open the Objective filter, select `Agentic`, then choose the specific objective to display only Agentic-mode findings. Check the Mode and Impact columns to confirm the attack mode and plain-language impact label (e.g. `Data disclosed: High`) for each row.

— Names exact columns, labels, and filter behavior with a usage path.

View the attack mode and impact of each…

03

Custom rule authoring across Endpoint Security policy resources

NEW

75

Endpoint Security Policies gain custom rule creation across three Policy Resources areas, all reachable via `Endpoint Security → Endpoint Security Configuration → Policy Resources`: custom Integrity Monitoring Rules that monitor registry values or files, or use custom XML (with support to duplicate, delete, and customize/reset TrendAI™ Integrity Monitoring rules); custom Intrusion Prevention Rules built from new custom application types (with support to configure/reset TrendAI™ Intrusion Prevention rules); and custom Log Inspection Rules authored in Basic format or custom XML targeting a specific log file path.

Create a custom rule to watch a sensitive registry value for unauthorized changes in an Endpoint Security Policy.

> 📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Policy Resources → Integrity Monitoring Rules, select `Create Custom Rule`, choose `Registry Value` as the monitor type, enter the target key and value name, then save and assign the rule to an Endpoint Security Policy.

Author a custom prevention rule to block traffic matching a known malicious pattern, assigned to the Advanced ruleset.

> 📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Policy Resources → Intrusion Prevention Rules, click `Create`, define your rule using a signature or start/end patterns, set the action to `Prevent`, and save. Assign the rule (Advanced ruleset) to the target Endpoint Security Policy.

When you need to inspect a custom application log file not covered by default rules, create a tailored Log Inspection rule scoped to that file.

> 📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Policy Resources → Log Inspection Rules, then select `Create` and choose Basic format or custom XML, specify the target log file path, and save the rule for use in an Endpoint Security Policy.

— Lists every rule type and surface but gives little mechanism beyond creation.

`Create custom Integrity Monitoring rule…` `Create custom Intrusion Prevention rule…` `Create custom Log Inspection rules in E…`

└──▷ BREAKING ON UPGRADE

! Smart Protection Services enrollment applies to ALL agents assigned a connection policy and cannot be undone — configure connection policies before enrolling.

WAS THIS USEFUL? 👍 👎

rustinel Feedback ⓘ v1.11.0 NOTES ↗ RANK

Rustinel is an endpoint detection for Windows, Linux, and macOS. Sigma, YARA, and IOC rules on native telemetry.

> Rustinel v1.11.0 adds agent tamper and lifecycle alerting, Windows code-signing verdicts on image and driver loads, broader Windows telemetry collection, wider Sigma rule coverage, and DNS event enrichment.

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

01 **Windows code-signing verdicts on image loads** NEW 85

Windows image and driver loads now expose `Signed`, `Signature`, and `SignatureStatus` fields with Sysmon semantics, covering embedded signatures, catalog-signed files, and AppX and MSIX package files. Image loads now also verify AppX and MSIX package files against their signed block map.

— Exact field names and verification mechanism given v1.11.0

02 **Expanded Windows telemetry fields and default events** NEW 80

Process starts now populate `LogonId` on both Windows and Linux, and `CurrentDirectory` on Windows. Windows kernel image loads are now classified as `driver_load` events. Security event 4740 is now collected by default on Windows, and Security 5379 joins the default subscription. New options `windows.security_domain_controller_events` (domain-controller events) and `windows.security_long_tail_events` (additional long-tail events) enable extra Windows Security event collection, both off by default.

— Config keys and event ids named verbatim v1.11.0

03 **Agent tamper detection and lifecycle alerting** NEW 78

A new `agent::tamper::etw_provider_modified` alert fires when an ETW provider is disabled while the sensor is running. Loss alerts now also cover ETW sessions, eBPF rings and maps, macOS ESF and BPF drops, and artifact queues. Alerts tagged `event.dataset: "edr.agent"` route lifecycle, tamper, and loss alerts to a separate pipeline that bypasses deduplication and active response.

— Names exact alert id and routing field, mechanism explained v1.11.0

04 **DNS event enrichment and eBPF capture** IMPROVED 65

DNS events gain derived fields including registered domain and name shape for use in Sigma rules, and Linux eBPF now captures DNS over TCP and `recvmmsg` responses

while flagging truncated (cut) names.

— Specific capture paths named, limited usage guidance v1.11.0

05 **Sigma rule coverage for antivirus and auditd** IMPROVED

60

Sigma rule support now extends to `category: antivirus` rules, which are routed to Defender detections, and `linux/auditd` rules, which are satisfied through an auditd field view.

— Names categories but little mechanism detail v1.11.0

THINNER COVERAGE BELOW

06 **SigmaHQ compatibility pin update** BREAKING

50

The SigmaHQ compatibility pin moves to `c04ad420`, so `sigma doctor` verdicts on the SigmaHQ rule tree may change after upgrade.

— Names exact pin and affected command, thin explanation v1.11.0

└──▷ **BREAKING ON UPGRADE**

! The SigmaHQ compatibility pin moves to `c04ad420`, so `sigma doctor` verdicts on the SigmaHQ rule tree may change after upgrade.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

Rapid7 InsightVM

Feedback ⓘ changelog · Nov 10, 2026 DOCS ↗ BLOG ↗

Rapid7 InsightVM is a vulnerability management platform that scans networks and assets to identify, prioritize, and help remediate security weaknesses.

NEW VULNERABILITY CHECK Authenticated detection for Node.js

NEW COMPLIANCE FRAMEWORK CIS macOS 26 Tahoe policy content

UPDATED SUPPORTED PLATFORM Oracle Database detection for symlinked homes

WAS THIS USEFUL? 👍 👎

defectdojo

Feedback ⓘ 3.5.100 NOTES ↗ RANK

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

> DefectDojo 3.5.100 adds a system-wide read-only mode for the platform.

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

01 **System-wide read-only mode** NEW 55

DefectDojo can be put into a system-wide read-only mode, configurable via `settings.dist.py` or `local_settings.py`.

— Names config files but not the exact setting key or scope 3.5.100

WAS THIS USEFUL? 👍 👎

OSV-Scanner

Feedback ⓘ v2.7.1 NOTES ↗

OSV-Scanner is a vulnerability scanner written in Go which uses the data provided by https://osv.dev

OSV-Scanner adds time-bounded vulnerability ignores via `EffectiveUntil`, warnings for stale `[[PackageOverrides]]`, a Fix available HTML report filter, and ecosystem-priority severity in table output.

└──▷ GET THIS VERSION ▾

Docker

```
$ docker pull ghcr.io/google/osv-scanner:v2.7.1
```

Go

```
$ go install github.com/google/osv-scanner/v2/cmd/osv-scanner@v2.7.1
```

git

```
$ git clone --branch v2.7.1 https://github.com/google/osv-scanner.git
# already have the repo? check out this version:
$ git checkout v2.7.1
```

published under the project's GitHub account · Go module path matches the repo · checked 10 Nov

Wrong command?

Warn about unused or expired `[[PackageOverrides]]` configuration entries so stale overrides do not silently persist.

Use `EffectiveUntil` in `[[IgnoredVulns]]` configuration entries to set expiry dates on vulnerability ignores; `IgnoreUntil` is now deprecated.

Add a Fix available filter checkbox to HTML reports.

Now shows ecosystem-priority severity (such as Ubuntu severity) in table output when available.

Enable the `os/ubuntu/binarytosource` enricher in the `vulns` preset to map Ubuntu binary package names from Chisel manifests to source package names.

+2 more

NEW SUPPORTED PLATFORM FreeBSD package extraction

NEW SUPPORTED PLATFORM Go workspace extraction

NEW SUPPORTED PLATFORM Haskell Cabal installed packages extraction

NEW SUPPORTED PLATFORM Python setup.cfg extraction

NEW SUPPORTED PLATFORM R installed packages extraction

NEW SUPPORTED PLATFORM Echo OS ecosystem mapping in DPKG extractor

UPDATED SUPPORTED PLATFORM Browser extension extraction now covers Edge and Firefox

UPDATED SUPPORTED PLATFORM PURL type mapping for additional ecosystems

└──▷ **BREAKING ON UPGRADE**

! The `haskell/cabal` extractor is renamed to `haskell/cabalprojectfreeze`; any configuration or scripts referencing it by name must be updated.

! The `chrome/extensions` extractor is renamed to `browser/extensions`; any configuration or scripts referencing it by name must be updated.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface Management

Censys

Feedback ⓘ November 9, 2026 (2026-11-09) BLOG ↗ RANK

Censys is a internet-wide scanning and search platform that catalogs hosts, services, and certificates to help security teams assess attack surface and exposure.

> Censys introduced a `PQC` label to flag post-quantum-cryptography-capable services and web properties, and significantly expanded its fingerprint database with over 200 new software, 40 new hardware, and 7 new OS fingerprints.

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

01 PQC label for post-quantum cryptography assets NEW

65

Censys now adds a `PQC` label to services and web properties on the Platform that support post-quantum cryptography, letting teams query and track PQC-capable assets across their attack surface. Users can filter asset search or attack surface views by the `PQC` label to list all identified services and web properties.

Find all internet-exposed services in your attack surface that support post-quantum cryptography using the new PQC label.

> 📍 Steps: In the Censys Platform, open the asset search or attack surface view and filter by label `PQC` to list all services and web properties identified as supporting post-quantum cryptography.

— Named label and filter usage given, but no query syntax or API shown

November 9, 2026 (2026-11-09)

02 Expanded software, hardware, and OS fingerprint coverage NEW

60

Censys adds over 200 new software fingerprints covering products such as Zabbix, CyberArk Privilege Cloud, CyberArk Privileged Access Manager Self-Hosted (PVWA), Palo Alto Networks Prisma Cloud Compute Console, FortiSOAR, Strapi 5, NGINX Plus, NGINX Unit, Cribl Stream, Apache Kafka, Apache Cassandra, etcd, Openfire XMPP Server, OPNsense, mySCADA myPRO HMI Runtime, Delta Electronics DIAEnergie, Delta Electronics DIAView, SAProuter, Rejetto HTTP File Server (HFS 2.x), and Salt (SaltStack) salt-api rest_cherrypy. It adds over 40 new hardware fingerprints for industrial control, networking, and OT devices including ABB AC500 V2 PLC CPU, Phoenix Contact PLCnext Control, Moxa EDR Series Secure Router, Hitachi Energy RTU500 Series RTU, Fortinet FortiAP Wireless Access Point, Cisco Catalyst series switches and access points, SMA inverter and controller products, Fronius Datamanager, Growatt ShineLink and ShineLan-X Datalogger, and Furuno FELCOM Ship Earth Station. It also adds 7 new OS fingerprints

including Sierra Wireless AirLink OS, Phoenix Contact AXC ProConOS eCLR, Veeam JeOS Appliance, Juniper Junos OS (SNMP), LANCOM LCOS WEBconfig, SMA ennexOS, and Victron Energy Venus OS.

— Extensive named products but no query or command shown for using them

November 9, 2026 (2026-11-09)

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

N-able N-central Feedback ⓘ

Introducing the New Windows OS Patching Engine (preview) in N-sight and N-central (2026-1

BLOG ↗

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

N-able N-central adds a new Windows OS patching engine to Endpoint Remediation, running independently of PME, with policies and additional controls coming soon.

NEW SUPPORTED PLATFORM Adds a new Windows OS patching engine in Endpoint Remediation for on-demand patching, running independently of the existing Patch Management Engine (PME) with no shared settings.

Both the new Windows OS engine and PME can run side by side without disabling either.

Upcoming policy support will cover automated approvals, schedules, reboot behaviour, and tag-based device targeting.

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ

What's New in Iru: Mac Privacy, SSO, and Managed Google Domain NOTES ↗ RANK

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

> Iru expanded its single release this window with Mac support for Privacy Library Items, an upgraded SSO sign-in experience, and in-product migration of Android devices to a managed Google domain.

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

01 **In-product Android managed Google domain migration** NEW 40

Iru now supports upgrading Android devices to a managed Google domain directly within the product.

— Names platform and action but lacks exact steps

What's New in Iru: Mac Privacy, SSO, an…

02 **Privacy Library Items extended to Mac** NEW 30

Iru now extends Privacy Library Item coverage to Mac computers, building on existing support for other platforms.



— Names the feature but no mechanism or scope detail

What's New in Iru: Mac Privacy, SSO, an…

03 **SSO sign-in experience upgrades** IMPROVED 20

Iru now supports Single Sign-On upgrades to improve the user sign-in experience.

— Vague wording, no specifics on what changed

What's New in Iru: Mac Privacy, SSO, an…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Ory

Feedback ⓘ changelog · Nov 10, 2026 CHANGELOG ↗ RANK

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

> Ory tightened security controls this window, enforcing stricter SAML assertion validation and locking down token lifespan configuration for dynamically registered OAuth clients.

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

01 SAML assertion issuer enforcement in Ory Polis

BREAKING

55

Ory Network (via Ory Polis) now accepts a SAML assertion only from the IdP of its connection, requiring the Issuer of a signed SAML assertion to match the entity ID of the configured IdP. Assertions from other issuers will be rejected.

– Explains mechanism and before/after but no config key given `changelog-20261110-1e3be56f`

02 Admin-only token lifespans for dynamic client registration

55

BREAKING

Ory Hydra now restricts token lifespan configuration to administrators for dynamically registered clients, removing per-client token lifespan overrides from OpenID Connect Dynamic Client Registration. OpenID Connect Dynamic Client Registration no longer accepts per-client token lifespan settings.

– Clear before/after but no specific field or flag named `changelog-20261110-1e3be56f`

└──▷ **BREAKING ON UPGRADE**

! OpenID Connect Dynamic Client Registration no longer accepts per-client token lifespan settings; token lifespans for dynamically registered clients are now administrator-only.

! Ory Polis now requires the Issuer of a signed SAML assertion to be the entity ID of the connection's IdP; assertions from other issuers will be rejected.

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ Auth for UCP is now in Early Access (2026-11-09)

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 shipped GA Dashboard Search with a global Command Palette and Boolean filtering, Early Access Management API search endpoints for Clients, Client Grants, Resource Servers and Organizations, and Early Access support for securing Universal Commerce Protocol (UCP) servers so AI agents can shop on users' behalf.

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

01 Dashboard Search reaches general availability IMPROVED

95

A global Command Palette (`Cmd+K` on macOS, `Ctrl+K` on Windows/Linux) now searches Applications, APIs, Organizations, and Users by name, display name, or ID from any Dashboard page. Guided filters and Boolean search are available in the search bar of Applications, APIs, Roles, Custom Domains, Organizations, Organization members, Enterprise Connections, and Database Connections, combining up to 5 filters with AND/OR logic, searching metadata by key and value or key only, and using NOT to exclude results; filter pills persist in the URL so filtered views can be shared or bookmarked.

Find a specific Organization from any Dashboard page without navigating to the Organizations section first.

> 📍 Press Cmd+K (macOS) or Ctrl+K (Windows/Linux) to open the Command Palette, then type the organization name, display name, or ID to jump directly to it.

Build a shareable filtered view of Enterprise Connections using multiple conditions and send it to a teammate.

> 📍 Navigate to Dashboard › Authentication › Enterprise Connections, use the search bar to add up to 5 AND/OR filter pills (e.g. by name or metadata key/value), then copy the URL — the filters are encoded in it and will restore when opened.

— Exact keystrokes, scope list, and URL persistence all given

Dashboard Search is now GA! (2026-11-09)

02 Management API search for Clients, Client Grants, Resource Servers, Organizations

NEW

85

New Early Access endpoints let you search directly instead of paging through full lists or filtering client-side: `GET /api/v2/clients/search`, `GET /api/v2/client-grants/search`, `GET /api/v2/resource-servers/search`, and `GET /api/v2/organizations/search`.

Find a specific client by name without iterating pages — useful in automation scripts that need to look up a client ID before provisioning.

```
$ curl -X GET 'https://<your-domain>/api/v2/clients/search?q=my-app-name' \
  -H 'Authorization: Bearer <mgmt-api-token>'
```

Locate an organization by identifier during an incident response workflow without client-side filtering across paginated results.

```
$ curl -X GET 'https://<your-domain>/api/v2/organizations/search?q=acme-corp' \
  -H 'Authorization: Bearer <mgmt-api-token>'
```

— Four exact endpoints plus runnable curl examples

Management API Search for Clients, Clie…

THINNER COVERAGE BELOW

03 **Auth for Universal Commerce Protocol (UCP)** NEW

45

Auth0 now supports securing a UCP server in Early Access so trusted AI agents such as Gemini can connect to it and act on behalf of shoppers, including discovering catalogs and making purchases. User identities can be linked so a shopper's identity in an AI agent connects to their account with a business, enabling personalization, loyalty points, and order tracking.

— Describes purpose and example agent but no setup steps

Auth for UCP is now in Early Access (20…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password SaaS Manager #2511 (2026-11-04)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password's biggest addition this window is Environments — shareable, per-project secret collections with matching support added across the CLI (`op environment secret list`), SDKs (list_secrets()), and the MCP server (`append_secrets` / `list_secrets`), alongside a new Events API token-introspection endpoint and new `environments` / `provisioning` CLI command groups.

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

01 **`op environment secret list` command** DEPRECATED 95

The CLI adds `op environment secret list { <environmentName> | <environmentID> | - }` to list secrets in a 1Password Environment, returning each secret's name, value, and concealment status; output can be formatted with `--format` (or `OP_FORMAT`) as `human-readable` , `json` , `csv` , or `dotenv` , and concealed values can be revealed with `--reveal` . Secrets are concealed in stdout/stderr by default; the CLI can also print a single named secret or run a script provisioned with Environment secrets directly. The prior command for listing Environment secrets is deprecated on November 17, 2026 and removed on January 20, 2027 — callers must migrate to `op environment secret list` .

Export all secrets from a staging environment into a dotenv file for local development, including concealed values.

```
$ op environment secret list Staging --reveal --format=dotenv > .env
```

Inspect all secrets in a specific environment by ID in CI, keeping concealed values hidden for safe audit logging.

```
$ op environment secret list 6ijsqfxvdxlwjcpsx7yrzs3tka --format=json
```

List all secrets in a specific 1Password Environment for a local development setup.

```
$ op environment secret list <environment-id>
```

— Full command, flags, env var and migration deadline given `product docs`

02 **MCP server tools for Environment secrets** NEW 77

The 1Password MCP server adds `append_secrets`, to add or update secrets in an Environment (recommended for rotating secrets previously stored in plaintext in code or `.env` files), and `list_secrets`, which returns only secret names, never values. Claude Code can import an existing env file into a new Environment via the 1password-environments skill — it creates the Environment, mounts it at the same file path, and stores values (API keys, tokens, passwords, private keys, connection strings) as concealed secrets, prompting to merge or rename if a duplicate Environment name is found. On macOS 27 `Golden Gate`, an `MCP error -32000: Connection closed` is fixed by enabling the `Data shared by 1Password` permission for the client app under Privacy & Security > Files & Folders.

Rotate a leaked secret by overwriting it in a 1Password Environment rather than leaving it in a plaintext .env file.

> 📍 "Add a secret for my OpenAI API key" (triggers `append_secrets` via your MCP client, e.g. Claude or Cursor)

Inspect which secrets are defined in a project environment before running a workflow that depends on them.

> 📍 "Show me the secret names in my project environment" (triggers `list_secrets` via your MCP client)

Troubleshoot MCP connection failures on macOS 27 Golden Gate when the 1Password server reports `MCP error -32000: Connection closed`.

> 📍 In System Settings, go to Privacy & Security › Files & Folders, find the Claude app (or your terminal/editor), and turn on the setting that starts with `Data shared by 1Password`. Quit and reopen the app, start a new Claude Code session, then check the server status.

Ask your MCP client to list what secrets exist in a project environment before adding new ones.

> 📍 Show me the secret names in my project environment

Rotate a secret that was previously stored in plaintext — tell your MCP client to store the new value in a 1Password Environment.

> 📍 Add a secret for my OpenAI API key

If your MCP client cannot connect after upgrading to macOS 27 Golden Gate, grant it file access to unblock the 1Password MCP server.

> 📍 Open System Settings › Privacy & Security › Files & Folders, select your MCP client app (e.g. Claude, Cursor, or Codex), and turn on the setting that starts with `Data shared by 1Password`. Then quit and reopen your MCP client.

— Named tools and exact fix path, but only natural-language triggers `product docs`

---

03 **Events API token introspection endpoint** NEW

75

The Events API adds a v2 token-introspection endpoint, `GET https://events.1password.com/api/v2/auth/introspect`, which inspects a bearer token's validity and metadata when called with an `Authorization: Bearer <token>` header.

Verify that a 1Password Events API bearer token is still valid before issuing requests against the Events API.

```
$ curl --request GET \
  --url https://events.1password.com/api/v2/auth/introspect \
  --header 'Authorization: Bearer <token>'
```

— Exact endpoint, method, header and runnable curl example `product docs`

---

04 **SDK `list_secrets()` function for Environments** DEPRECATED

63

1Password SDKs add a list_secrets() function to read secrets from an Environment directly — it accepts an Environment name or ID plus a `reveal` option and returns a list of `EnvironmentSecret` objects without the previous response wrapper. The existing secrets-reading function and the `EnvironmentVariable` types it returned are deprecated on November 17, 2026 and removed on January 20, 2027; callers must migrate to list_secrets().

— Function and types named but no runnable code example given `product docs`

---

05 **Environments for organizing and sharing secrets** NEW

60

1Password introduces Environments — shareable, context-scoped secret collections for each project, application, or deployment stage (e.g. staging or production), replacing per-project `.env` file setup. Team members on shared accounts can be granted per-Environment access to keep secrets for different projects or stages isolated, and secrets on an Environment's details page can be searched or sorted alphabetically via the `Search secrets` field. Secrets can be added manually (desktop app `New secret`) or imported in bulk from an existing `.env` file, and secrets loaded from an Environment

can be combined with secret references from a `.env` file in the same invocation via 1Password CLI or SDKs.

Set up a development environment's secrets and load them into a local script without writing plaintext secrets to disk.

> 📍 Steps: 1. Open and unlock the 1Password desktop app. 2. Navigate to your project's Environment. 3. Select `+ Add secret`, enter the secret name in the Name field and its value in the Value field. 4. Repeat for each secret. 5. Use 1Password CLI or an SDK to read and pass the secrets to your application as environment variables.

Add a secret manually to an Environment during local project setup when you don't have an existing .env file to import.

> 📍 In the 1Password desktop app: open and unlock the app, navigate to your Environment, select `New secret`, enter the secret name in the Name field and its value in the Value field, then repeat for each secret.

Bootstrap a project environment quickly by importing all variables from an existing .env file instead of adding secrets one by one.

> 📍 In the 1Password desktop app: open and unlock the app, navigate to your Environment, then select the option to import a file — 1Password will automatically import the environment variables from the file as secrets.

— UI and conceptual workflow described, no CLI/API flags named here `product docs`

THINNER COVERAGE BELOW

06 New `environments` and `provisioning` CLI command groups NEW

33

The CLI reference adds two new top-level command groups: `environments`, for managing 1Password Environments and their secrets, and `provisioning`, for managing provisioning integration settings.

— Only names two command groups, no subcommands detailed `product docs`

UPDATED INTEGRATION Atlassian Cloud and Opsgenie account consolidation

UPDATED INTEGRATION Cloudflare Access log timeout retry

UPDATED INTEGRATION Zenefits employee sub-type values in custom person-type rules

UPDATED INTEGRATION Contentful authorization scope fix

└──▷ BREAKING ON UPGRADE

! The existing secrets-reading function and the response-wrapper types it returns will be deprecated on November 17, 2026 and removed on January 20, 2027; callers must migrate to list_secrets().

! The existing command to list secrets in a 1Password Environment via 1Password CLI will be deprecated on November 17, 2026 and removed on January 20, 2027 — replace it with `op environment secret list`.

! The existing function in 1Password SDKs and the `EnvironmentVariable` types it returns will be deprecated on November 17, 2026 and removed on January 20, 2027 — replace it with list_secrets().

> `1Password SaaS Manager #2511 (2026-11-04)` BLOG ↗ LATEST STABLE
>
> 1Password SaaS Manager adds scoped budgets and spend limits, groups OpenAI usage records, adds human-readable Claude API key names, and expands log redaction.
>
> Create scoped budgets and spend limits for applications, teams, cost centers, and API keys.
>
> OpenAI request-level usage is now grouped into more useful records, reducing noise from individual request identifiers.
>
> Claude usage imports now include human-readable API key names.
>
> Jira-connected tasks can now be closed or deleted manually when the related Jira issue is closed or deleted.
>
> Log redaction now covers additional secrets and preserves boolean values, and reduces sensitive provider, user, and vendor response data written to integration and workflow logs.

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

Aembit

Feedback ⓘ

Aembit Now Secures Google Antigravity Agents (2026-11-09) BLOG ↗ RANK

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

> Aembit's centerpiece this window is MCP Identity Gateway, a managed (or self-hosted) proxy that enforces Access Policies and injects per-request credentials between AI agents and MCP servers, paired with blended user+agent identity policies, new MCP access reporting, Microsoft Copilot Studio support, SSO/OIDC-group controls on Client Workload identification, and a Windows 11 Aembit Desktop with a credential-free local proxy.

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

01 MCP Identity Gateway for AI agent-to-MCP-server access NEW

98

Aembit adds MCP Identity Gateway, a stateless data-plane proxy reachable at `https://<tenant-id>.mcpgateway.aembit.io` that enforces Access Policies on every AI-agent-to-MCP-server request and injects credentials per request so agents never hold downstream secrets; it can also be self-hosted on a Linux host you provision, configured via environment variables Aembit sets during provisioning, keeping all runtime state within one Aembit Tenant. The gateway validates a signed JWT against Aembit's JWKS signing keys on every request, caches downstream credentials per-request for 60 seconds without persisting them to disk, denies all requests by default (fail-closed) unless a policy explicitly allows them, and produces structured logs of agent identity, user identity, target server, and policy decisions. Access is governed by two Access Policies — a left-side Client-to-Gateway policy and a right-side Gateway-to-Server policy — with optional time-based and GeoIP-based Access Conditions and a Content Security provider enforcing four boundaries (token validation against the Trust Provider, TLS termination/re-initiation, in-memory-only credential handling, and per-request policy evaluation); the gateway auto-discovers available MCP servers, proxies resource discovery/retrieval through the same flow as tools, and provides per-user credential isolation with centralized routing to multiple MCP servers from one endpoint.

Verify your tenant's MCP Identity Gateway endpoint is live and TLS is terminating correctly before configuring any client workloads.

```
$ curl -v https://<tenant-id>.mcpgateway.aembit.io/
```

Walk through the two required Access Policies and optional Content Security provider in the Aembit console to fully wire up an AI agent to an MCP server.

> 📍 1. In the Aembit console, go to Access Policies and create the Client-to-Gateway policy: set the Client Workload to the MCP client, set the Server Workload hostname to <tenant-id>.`mcpgateway.aembit.io`, and attach a Credential Provider that issues tokens whose audience matches MCP Identity Gateway. 2. Create the Gateway-to-Server policy: set the Client Workload to MCP Identity Gateway (https://<tenant-id>.aembit.io), set the Server Workload to the upstream MCP server, and configure OAuth with the Authorization URL and Token URL from the MCP server vendor. 3. Optionally, open the Access Conditions card in the right panel on either policy and add time-based or GeoIP-based conditions. 4. Optionally, attach a Content Security provider to the Gateway-to-Server policy to inspect each MCP tool message.

Route an AI agent through MCP Identity Gateway instead of directly to an MCP server, so the agent authenticates through Aembit and holds no downstream credentials.

> 📍 In your Aembit Tenant, go to MCP Identity Gateway › Access Policies and create a Gateway-to-Server policy targeting the MCP server you want to front. Point your AI agent at the MCP Identity Gateway endpoint in place of the direct MCP server URL.

— Mechanism, defaults and architecture fully detailed with runnable examples.

product docs

---

02 **Aembit Desktop for Windows 11 with credential-free proxy** NEW

93

Aembit Desktop is a new Windows 11 x64 (public beta) application for accessing HTTP and HTTPS Server Workloads without static credentials, via a local proxy that injects service credentials into requests per-policy with no credentials stored on the machine. Administrators deliver managed Desktop settings as JSON files to override per-user configuration, and users can sign in through an existing SSO Identity Provider once Enable User Access is toggled on that Identity Provider's Users tab (visible only after Aembit support enables Desktop for the tenant). The `adt beta tool aws` command runs AWS CLI commands without storing AWS access keys on the machine, pulling credentials from an Access Policy at request time, e.g. `adt beta tool aws s3 ls s3://my-bucket`.

Let Windows 11 users sign in to Aembit Desktop through an existing SSO Identity Provider without static credentials.

> 📍 In the Aembit console, go to Administration › Identity Providers › <your IdP> › Users tab, then toggle Enable User Access and save. The Users tab appears only after Aembit support enables Aembit Desktop for your tenant.

Run an AWS CLI command from a developer workstation without any stored AWS access keys — Aembit Desktop injects the credential from the matching Access Policy.

```
$ adt beta tool aws s3 ls s3://my-bucket
```

Run an AWS CLI command from a developer workstation without storing AWS access keys — the credential comes from the Access Policy at request time.

```
$ adt beta tool aws s3 ls s3://my-bucket
```

— Exact command, toggle path and file-based settings mechanism given. product docs

03 **Microsoft Copilot Studio agent support in MCP Identity Gateway**

NEW 90

Connect Microsoft Copilot Studio agents to MCP Identity Gateway over Streamable HTTP transport, with the gateway validating OAuth 2.0 access tokens and proxying requests to downstream MCP servers using Aembit-managed credentials. For servers without OAuth discovery, such as Microsoft Graph via the Microsoft MCP Server for Enterprise, configure the callback URL manually as `https://<tenant-id>.id.<stack>.aembit.io/mcp-auth/userauth/<cp-id>/callback`. To stay within Copilot Studio's 70-tool-per-agent limit, assign fewer Access Policies per agent, use per-tool toggles, split integrations across child agents with their own tool budgets, or keep MCP tool names to 64 characters or fewer including the Server Workload prefix MCP Identity Gateway adds.

Connect a Copilot Studio agent to MCP Identity Gateway when the target MCP server (such as Microsoft Graph) does not support OAuth discovery and requires manual OAuth configuration.

📍 In Copilot Studio, create a tool with your MCP Identity Gateway endpoint (ending in `/mcp`) as the Server URL, then when prompted for OAuth values enter the callback URL manually: https://<tenant-id>.id.<stack>.aembit.io/mcp-auth/userauth/<cp-id>`/callback`

— Exact callback URL format and concrete workaround limits given. product docs

04 **SSO enforcement and OIDC group matching for Client Workloads**

NEW 90

Client Workloads gain an Enforce SSO setting (available on both Redirect URI and CIMD Client ID identification) requiring users to authenticate through a configured SSO

identity provider during the OAuth flow — a multi-select dropdown lists active SSO providers and at least one must stay selected or the save is rejected; for CIMD Client ID workloads the toggle also limits the MCP sign-in page to only the assigned providers. A single Client Workload can now carry both a Redirect URI and a CIMD Client ID identifier at once. A new `User Group` identifier lets Access Policies match MCP Identity Gateway users by the groups they belong to in an OIDC Identity Provider (via `User Groups & MCP Policy Access` group-lookup settings on the provider), and Aembit now supports automatic user creation tied to Identity Providers at sign-in. Separately, the MCP Authorization Server now accepts an `application_type` of `native` or `web` via Dynamic Client Registration, allowing native clients to register custom-scheme redirect URIs such as `vscode://`, and MCP Server Workload `URL Path` values can now end in a trailing slash, such as `/mcp/`.

Enforce that all Client Workload OAuth flows go through a specific SSO identity provider, preventing unauthenticated or direct access.

> 📍 In the Aembit console, go to Access Policies › Client Workloads › [your workload] › Identification › Redirect URI, enable `Enforce SSO`, then select at least one active SSO identity provider from the multi-select dropdown and save.

Disable SSO enforcement for a Client Workload to allow the OAuth flow to complete without user authentication — useful for fully automated, non-human workloads.

> 📍 In the Aembit console, open the Client Workload, go to MCP Authorization Configuration, and set Enforce SSO to `Off`.

Restrict which identity providers users can choose on the MCP authentication page for a specific Client Workload.

> 📍 In the Aembit console, open the Client Workload, go to MCP Authorization Configuration, set Enforce SSO to `On`, then deselect any identity providers you do not want shown on the MCP authentication page.

Let Windows 11 users sign in to Aembit Desktop through an existing SSO Identity Provider without static credentials.

> 📍 In the Aembit console, go to Administration › Identity Providers › <your IdP> › Users tab, then toggle `Enable User Access` and save. The Users tab appears only after Aembit support enables Aembit Desktop for your tenant.

Restrict an MCP server to users who belong to the `engineering` team or any group starting with `finance-` in your OIDC Identity Provider.

> 📍 In the Aembit console, open the Client Workload, add a `User Group` identifier, select the OIDC 1.0 Identity Provider, and enter `engineering` or `finance-*` as the

group value. Then go to the Identity Provider's Settings and enable `User Groups & MCP Policy Access`, setting the group claim name and cache time.

Restrict an MCP server to members of a specific engineering group defined in an OIDC Identity Provider, using wildcard patterns for teams like all finance sub-groups.

📍 In the Client Workload editor, add a `User Group` identifier, select your OIDC 1.0 Identity Provider, and enter `engineering` or `finance-*` as the group pattern. Enable `User Groups & MCP Policy Access` on the Identity Provider's settings to set the group claim and cache time.

Enforce SSO for an MCP client identified only by CIMD Client ID, so only users from approved Identity Providers can reach the MCP server.

📍 In the Client Workload editor, set the identifier to a CIMD Client ID, toggle on `Enforce SSO`, and select one or more Identity Providers. Save requires at least one Identity Provider selected.

— Every toggle, identifier and config value named with exact steps. `product docs`

---

05 **MCP access reporting and live authorization view** NEW

75

Aembit surfaces `MCP Access Activities` reports, listing MCP requests with no error or exception code grouped by user, MCP client, MCP server, method, and target via the Reporting dashboard; `MCP Access Exceptions` reports list requests carrying an error or exception code grouped the same way plus exception type. A live authorization request view at MCP Identity Gateway surfaces the redirect URI, resource, matched Client Workload, and Access Policy outcome for each inbound request ahead of the standard event-pipeline delay, and every MCP request is logged with both verified user identity and AI agent workload identity, capturing authentication events (user, agent, time, location) and workload events (resources accessed, operations performed, credentials issued) for compliance audit trails.

Review failed or blocked MCP requests at a glance — go to the Reporting dashboard and open MCP Access Exceptions to see every request that returned an error or exception code, grouped by user, client, server, method, target, and exception type.

📍 In the Aembit Tenant, go to Reporting › MCP reports › MCP Access Exceptions to view all MCP requests with an error or exception code, grouped by user, MCP client, MCP server, method, target, and exception.

Review failed or blocked MCP requests in one place to triage access policy issues or client misconfigurations.

📍 In the Aembit console, go to Reporting › MCP Access Exceptions. Open any row to see the exception category and code, the Access Policy decisions applied, and the per-server outcome for that request.

— Named dashboards and fields with a clear navigation path. product docs

06 **Blended identity combining user and agent identity** NEW

70

Access policies can now evaluate blended identity — combining user identity from your IdP with AI agent workload identity — so a single policy can ask whether this user, using this AI agent, is authorized to reach a resource right now, enabling per-agent-per-user rules such as allowing Confluence access through Claude Desktop but denying it through an untrusted third-party agent. The mechanism spans two Access Policies: a Credential Provider embeds the user's identity into the access token via dynamic claim injection, and MCP Identity Gateway uses its workload identity plus the embedded user context to obtain downstream credentials, storing OAuth tokens per user rather than per agent so revoking one user's access doesn't affect others sharing the same gateway.

— Mechanism and scope explained, but no example command given. product docs

THINNER COVERAGE BELOW

07 **MCP Server for AI-assistant tenant management** NEW

30

Aembit adds an MCP Server that exposes Aembit Tenant configuration to MCP clients, so practitioners can query and manage Aembit from an AI assistant.

— Only a one-line description, no interface or command given. product docs

NEW INTEGRATION Google Antigravity agent support

Aembit Now Secures Google Antigravity Agents (2026-11-09) BLOG ↗ LATEST STABLE

Aembit adds Google Antigravity support via the MCP Identity Gateway, brokering short-lived credentials using a blended agent-and-user identity model with full audit logging.

Point `mcp_config.json` at the Aembit MCP Identity Gateway to enforce real-time policy evaluation for Google Antigravity (`agy`) agents before any credential is issued.

Aembit's blended identity model evaluates both the agent's identity and the human user's IdP identity together, replacing long-lived Google ADC sessions and bearer tokens stored in `mcp_config.json` or cached at

> `~/.gemini/antigravity/mcp_oauth_tokens.json` with short-lived, scoped credentials.
>
> Every credential issuance and access event is logged against both the agent identity and the human identity behind it, making 'which agent touched production, and on whose behalf' a query rather than an investigation.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

Panorays

Feedback ⓘ updated · Nov 10, 2026 NOTES ↗ RANK

Panorays, a Forrester Wave Leader, helps 500+ companies assess, monitor and fix third-party cyber risk based on business context.

> Panorays rolled out an MCP server integration under Panorays Axis for AI/LLM-driven third-party risk workflows, broadened its public API around the questionnaire lifecycle, and added new widgets and filters to Live Dashboards.

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

01 **MCP server integration via Panorays Axis** NEW 55

Panorays Axis now connects AI/LLM tools to Panorays tenant data through a remote MCP server, with guided workflows for questionnaire analysis, assessment reports, and supplier onboarding. API documentation adds a dedicated section covering Panorays Axis (MCP).

— Describes mechanism and use cases but no exact endpoint or command given

snapshot-20261110

02 **Public API expansion for questionnaire lifecycle** IMPROVED 45

The public API now supports questionnaire lifecycle management, evidence validation retrieval, and expanded questionnaire data fields, with documentation adding dedicated sections for TPCRM lifecycle use cases and additional endpoint references under the Panorays domain.

— No specific endpoint paths, methods, or field names provided snapshot-20261110

03 **New widgets and filters in Live Dashboards** IMPROVED 35

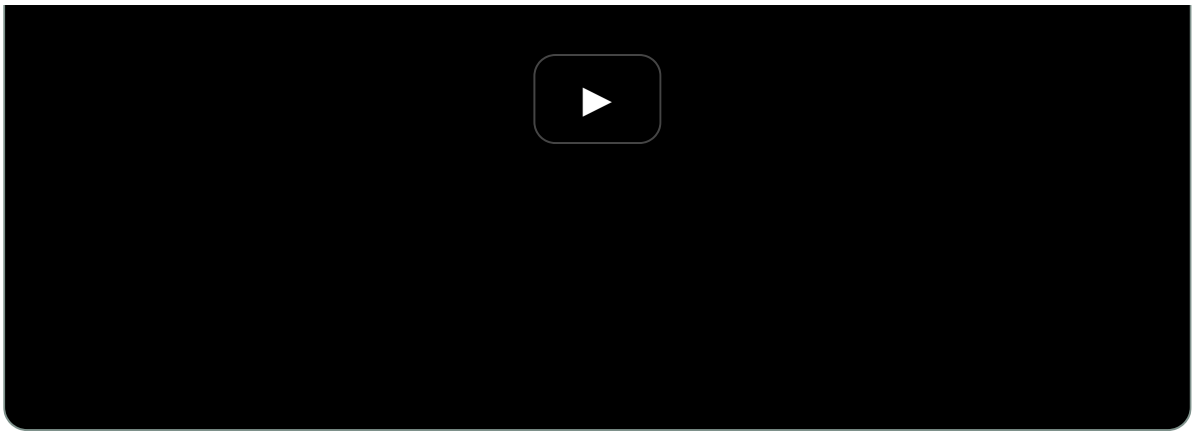
Live Dashboards now include new interactive widgets, pre-configured dashboards, and expanded filtering options across Panorays data points.

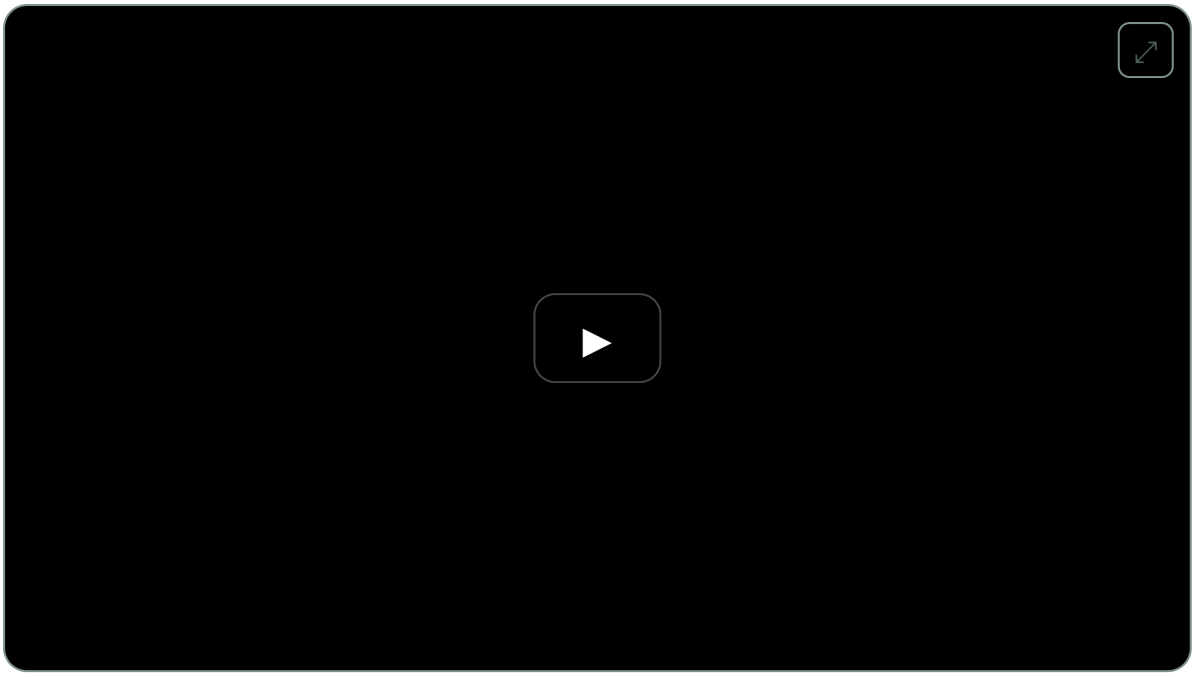
— No specific widget names, dashboard templates, or filter fields listed

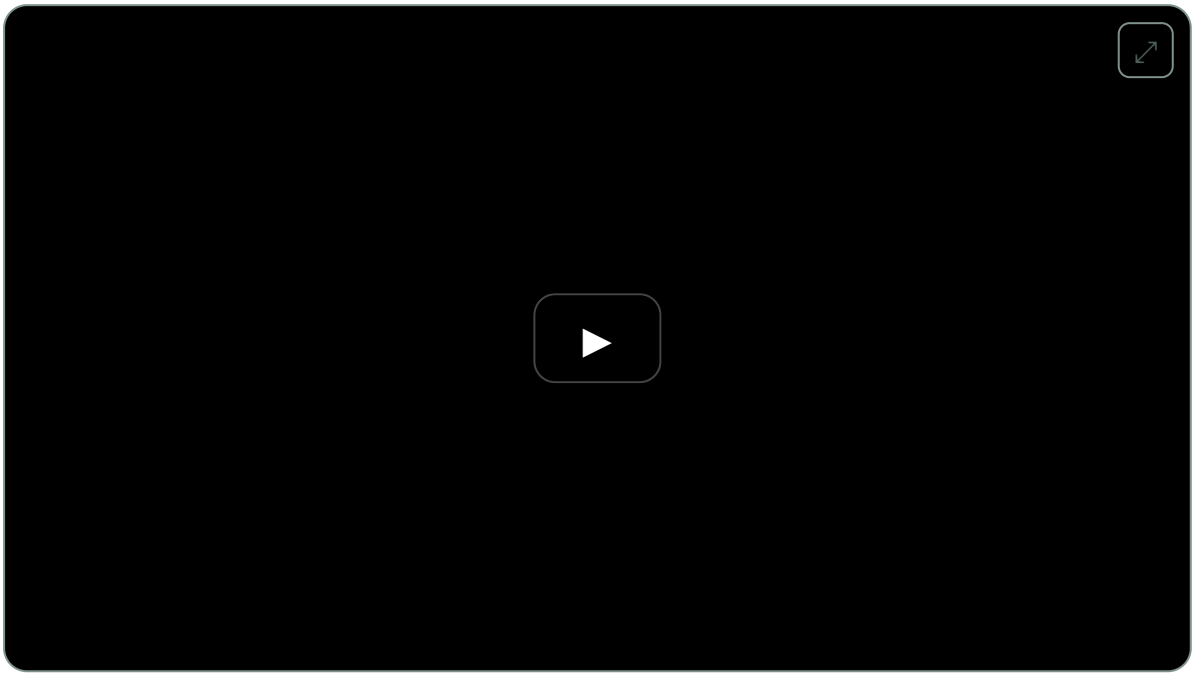
snapshot-20261110

└──▷ **ALSO FROM THESE RELEASES**









WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 4.0.0 (2026-11-09) BLOG ↗ RANK

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

> Runlayer's 4.0.0 release adds LLM gateway multi-model failover with stricter routing-rule validation, new self-hosted secret management options, and a broad set of MCP-driven controls spanning agent administration, webhook configuration, poll triggers, and eval workflows.

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

01 LLM gateway multi-model failover with routing-rule validation

95

BREAKING

LLM gateway routing rules can now set up to three fallback models so the gateway retries the next fallback after a server error, rate limit, timeout, or refused key, across providers if needed, configured via Settings → AI models → Budgets & routing. Creating or updating a routing rule via the API or Terraform now returns a `422` error if the rule's target or any fallback isn't listed under your AI providers; add the model under Settings → AI models → Providers or pull a valid id from `GET /api/v1/llm-gateway/models`. Existing rules keep routing, and updates that leave targets unchanged still succeed.

Discover which LLM models are available to use as routing-rule targets or fallbacks before creating a rule.

```
$ curl -s -H 'Authorization: Bearer <token>' \
  https://your-runlayer-instance/api/v1/llm-gateway/models | jq
'.[].id'
```

— Names endpoint, error code, UI path and fallback limit. 4.0.0 (2026-11-09)

02 Provisioned secret source for ECS deployments NEW

81

Self-hosted ECS deployments can set `tenant_secret_source = "provisioned"` so WorkOS, Sentry, and Resend keys are read directly from your AWS Secrets Manager secret instead of passing through Terraform state.

Keep WorkOS, Sentry, and Resend secrets in AWS Secrets Manager only, preventing them from appearing in Terraform state on ECS deployments.

```hcl
tenant_secret_source = "provisioned"
```

— Exact config key and value plus Terraform example. 4.0.0 (2026-11-09)

03 **RESEND_API_KEY for self-hosted email delivery** NEW

78

Self-hosted deployments can supply the Resend API key as the `RESEND_API_KEY` environment variable or the `runlayer-app` Kubernetes secret key, which takes priority over the WorkOS Vault secret fallback.

Inject the Resend API key into a Kubernetes self-hosted deployment so email delivery does not depend on WorkOS Vault.

```
$ kubectl create secret generic runlayer-app \
  --from-literal=RESEND_API_KEY=re_live_xxxxxxxxxxx \
  --dry-run=client -o yaml | kubectl apply -f -
```

— Exact env var, secret name, and kubectl command given. 4.0.0 (2026-11-09)

04 **Webhook signing and payload editing via MCP** NEW 75

Through the Runlayer Assistant or an MCP client, users can configure request signing, rotate or disable signing secrets, and rename agent webhooks. The webhook JSON body can also be edited in Settings → Policies → Approval settings using variables like `{{request_id}}`, `{{session_id}}`, and `{{status}}`, with a Send test option to try the payload.

— Names template variables and UI path for the editor. 4.0.0 (2026-11-09)

THINNER COVERAGE BELOW

05 **Agent administration controls via MCP** NEW 55

Through Runlayer MCP or the Runlayer Assistant, users can run a saved agent schedule once on demand without changing its next or last run time, transfer agent ownership to another person (subject to the same permission checks as the agent's Share dialog), and generate an AI icon suggestion for an agent that can be applied with `update_agent`.

— Names MCP tool and dialog but no full workflow detail. 4.0.0 (2026-11-09)

06 **Poll trigger introspection via MCP** NEW 53

Through Runlayer MCP, users can list a poll trigger's recent polls, run Fetch now, and list live choices for a trigger's target fields, such as Slack channels or GitHub repositories.

— Names concrete MCP actions but lacks mechanism detail. 4.0.0 (2026-11-09)

07 **Eval workflow management via MCP** NEW 53

Through the Runlayer Assistant or MCP clients, users can delete an eval, record a human verdict on an eval case, pick source runs from a plain-language description, label runs pass or fail, and group failed runs by cause.

– Lists actions but no command or endpoint named. 4.0.0 (2026-11-09)

└──▷ **BREAKING ON UPGRADE**

! Creating or updating an LLM gateway routing rule via the API or Terraform now returns a 422 error when the rule's target or any fallback is not a model listed on your AI providers. Add the model under Settings → AI models → Providers or use a model from `GET /api/v1/llm-gateway/models`. Existing rules continue routing; updates that leave targets unchanged still succeed.

WAS THIS USEFUL? 👍 👎

◆ **Agent Sandboxes**

microsandbox

Feedback ⓘ v0.7.14 NOTES ↗ RANK

Microsandbox is an easy, fast, programmable and local-first microVM runtime and library

> Microsandbox adds inbound TCP connection queuing and admission metrics for sandbox network traffic.

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

01 **Inbound TCP queuing with admission metrics** NEW 30

microsandbox now queues inbound TCP connections and exposes admission metrics for sandbox network traffic.

– No specifics on metric names or config to use them v0.7.14

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

Reality Defender

Feedback ⓘ changelog · Nov 10, 2026

CHANGELOG ↗ RANK

Reality Defender detects synthetic audio, video, images, and documents for enterprise deepfake screening.

> Reality Defender tightened scan timeout handling and free-tier enforcement with explicit error codes, improved audio analysis accuracy, and expanded YouTube URL ingestion to cover non-MP4 files and Shorts links.

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

01 **Stalled scan timeout reduced to 2 hours** BREAKING 90

Scans stuck in `ANALYZING` now automatically end after 2 hours instead of the previous 12 hours, with `status` on `GET /api/media/users/{requestId}` set to `UNABLE_TO_EVALUATE`.

Poll a scan and handle the new 2-hour timeout: check whether `status` is `UNABLE_TO_EVALUATE` and resubmit if so.

```
$ curl -s -H "Authorization: Bearer <token>"
https://api.realitydefender.com/api/media/users/<requestId> | jq
'.status'
```

— Named endpoint, exact status values, before/after timing, runnable example

changelog-20261110-8bb9aaef

02 **Free API plan blocks audio uploads with explicit error** 75

BREAKING

Audio uploads on the Free API plan now return HTTP 403 with code `free-tier-restriction` and message 'Forbidden: Video, audio, and text uploads require a paid plan', instead of failing silently as before.

— Exact HTTP code and error message given, but no sample call changelog-20261110-8bb9aaef

THINNER COVERAGE BELOW

03 **YouTube URL ingestion improvements** IMPROVED 55

YouTube URLs pointing to non-MP4 files are now downloaded and analyzed instead of failing, with failed downloads ending promptly rather than staying in `DOWNLOADING`;

YouTube Shorts links submitted to the upload endpoint are now downloaded and analyzed as well.

— Names status state and two ingestion cases, no endpoint or command given

changelog-20261110-8bb9aaef

---

04 **Audio analysis now detects previously-missed speech** IMPROVED

35

Audio analysis now covers speech that was previously mistaken for silence and excluded from results; scores for existing files may differ as a result of this change.

— Describes effect but no mechanism, numbers, or surface named

changelog-20261110-8bb9aaef

└──▷ **BREAKING ON UPGRADE**

! Audio uploads on the Free API plan now return HTTP 403 (`free-tier-restriction`) where they previously may have succeeded — upgrade to a paid plan to restore access.

! Scans that remain in `ANALYZING` for 2 hours now resolve to `UNABLE_TO_EVALUATE` ; integrations that expected up to 12 hours before timing out must be updated.

WAS THIS USEFUL? 👍 👎

Knostic

Feedback ⓘ

More Than Markdown: Hidden Data Collection in Two MCP Packages (2026-11-09)

BLOG ↗ RANK

Knostic enforces need-to-know access controls for enterprise AI assistants such as Microsoft Copilot.

> Knostic's AgentMesh scanner added detection rules flagging two malicious MCP packages that secretly exfiltrate secrets at server launch.

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

01 AgentMesh detection of mmdd-doc-mcp exfiltration NEW

75

AgentMesh now rates `@ifasn/mmdd-doc-mcp@2.1.1` as Dangerous, detecting that it exfiltrates process command-line secrets over plain HTTP to `121[.]199[.]24[.]94` via the request path `/beacon?stage=deepaudit&src=mmd-doc-mcp&res=` on each server launch.

— Names exact package, rating, IP and path; no remediation step given

More Than Markdown: Hidden Data Collect…

02 AgentMesh detection of ms-doc-mcp environment harvesting NEW

75

AgentMesh now rates `@ifasn/ms-doc-mcp@2.1.1` as Dangerous, detecting that it records full values of environment variables matching `KEY`, `TOKEN`, `SECRET`, `ALIBABA`, `DASHSCOPE`, or `CREDENTIAL` to `output/env_audit.txt` and stderr as `ENV-VALUE:` lines on each server launch.

— Names exact package, env-var patterns and output file; no remediation step given

More Than Markdown: Hidden Data Collect…

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Arcade

Feedback ⓘ changelog · Nov 10, 2026 CHANGELOG ↗ RANK

Arcade is the actions runtime for enterprise AI agents â one control point to enforce, execute, and govern every action, every time.

> Arcade shipped a batch tool-execution history deletion command, paged Google Drive search, new GitHub, Linear, Slack, and Teams tools, per-execution API key and gateway reporting, and a breaking overhaul of file-conflict handling across its Office and Figma toolkits.

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

01 Linear toolkit tool and field enhancements IMPROVED

80

`ListIssues` can now filter results by `creator` (a specific person or the caller) and returns completion time, cancellation time, creator, and a sub-issue's parent; `GetIssue` returns completion and cancellation time; you can clear an issue's parent or project, or set a project's priority; `ListWorkflowStates` returns each workflow state's team key alongside its ID and name; and a new `ListRecentComments` tool lists every comment created or edited across a team or project since a given time, including quoted inline comments.

— Names every tool and field touched, but no code sample changelog-20261110-6dc346e7

02 File-conflict handling in Excel and Word BREAKING

80

Excel upload and create calls replace the previous plain `overwrite` flag with `conflict_behavior` (`fail`, `rename`, or `replace`); Word's create call now refuses a same-named file by default instead of silently replacing it, controlled by the same `conflict_behavior` argument. Existing calls using the old overwrite flag require migration.

Avoid accidental overwrites when creating a Word document by explicitly choosing what happens on a name conflict.

```yaml
conflict_behavior: rename
```

— Exact config values and a before/after, with example changelog-20261110-6dc346e7

03 GitHub search and comment-listing tools NEW

70

New GitHub tools `QueryIssues` and `QueryPullRequests` search issues and pull requests across repositories using GitHub's search syntax, and `ListRepositoryConversationComments` lists every issue and pull request conversation comment in a repository in one paged, filterable and sortable call.

— Names the three new tools and their call behavior changelog-20261110-6dc346e7

04 **OpenFGA database creation control** IMPROVED 70

Setting `rbac.openfga.database.create: false` makes the engine skip running `createdb` against the `postgres` maintenance database, letting operators point OpenFGA at a database a DBA already created with a least-privilege role.

Let OpenFGA use a pre-created, least-privilege database instead of trying to run createdb itself.

```yaml
rbac:
  openfga:
    database:
      create: false
```

— Exact config key, default behavior, and example provided changelog-20261110-6dc346e7

05 **Batch tool execution history deletion command** NEW 60

A new engine command deletes an organization's stored tool execution history on demand; it defaults to a dry run, requires a reason, and caps the number of rows removed per run.

— Mechanism described in detail but exact command name missing

changelog-20261110-6dc346e7

06 **Slack bulk file deletion tool** NEW 60

The new Slack `DeleteFiles` tool deletes up to 10 files by ID or link in one call, reporting a success, rejection, or unknown outcome per file.

— Names tool, limit, and per-file outcomes changelog-20261110-6dc346e7

THINNER COVERAGE BELOW

07 **Paged Google Drive file search** IMPROVED 50

Google Drive's `SearchFiles` can now page results past the first page using a continuation token.

— Names tool and mechanism but no code sample changelog-20261110-6dc346e7

08 **Teams meeting artifacts listing tool** NEW 50

The new Teams `ListMeetingArtifacts` tool lists a meeting's recordings and transcripts in one call, reporting when an artifact is not available.

— Names tool and behavior, no further mechanism changelog-20261110-6dc346e7

09 **Figma toolkit tools removed** BREAKING 40

The Figma toolkit's `ListProjects` and `ListFiles` tools are removed, requiring operator action on upgrade.

— Names removed tools but no migration path given changelog-20261110-6dc346e7

10 **Per-execution API key and gateway reporting** NEW 30

Arcade now records which API key and gateway each tool execution ran through, reported alongside the execution's other details.

— States what's recorded but no access path given changelog-20261110-6dc346e7

11 **Broader file scopes in Office toolkits** BREAKING 25

Microsoft Word, Excel, and PowerPoint toolkits move to broader file scopes, requiring operator action on upgrade.

— No specifics on what the new scopes are changelog-20261110-6dc346e7

└──▷ BREAKING ON UPGRADE

! Excel upload and create calls: the `overwrite` flag is replaced by `conflict_behavior` (`fail` , `rename` , `replace`) — existing calls using the old flag require migration.

! Microsoft Word, Excel, and PowerPoint toolkits move to broader file scopes — operator action required on upgrade.

! Word's create call now refuses a same-named file by default (instead of silently replacing it); the new `conflict_behavior` argument controls the behavior.

! Figma toolkit: `ListProjects` and `ListFiles` tools are removed — operator action required on upgrade.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 10, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue