THE CYBER TOOLCHAIN

NO. 131

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 5, 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 139 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 44 of 45 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 | Opus 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 today: JumpCloud can now stream identity events into SIEMs and export access reviews; Upwind turns

vim ? help

vulnerability findings into native ServiceNow Vulnerable Items; StepSecurity inventories environment secrets and checks their rotation, usage, and OIDC posture.

GOVERN

Turn access reviews into reusable audit evidence

Compliance-grade Access Certification exports move review results out of the console in a form teams can hand to auditors or retain with other evidence. The same release adds continuous SIEM ingestion of identity events, reducing the gap between identity administration and security monitoring.

JumpCloud

PROTECT

Route cloud vulnerabilities into the workflow built for fixing them

Upwind now files findings in ServiceNow Vulnerability Response as native Vulnerable Items instead of generic ITSM tickets. That preserves vulnerability-specific fields and lifecycle handling, so remediation owners can work from the system they already use for vuln response rather than translating tickets by hand.

Upwind

PROTECT

Find and govern GitHub environment secrets, not just repository secrets

StepSecurity added environment secrets to its Action Secrets page, including inventory, rotation, usage, and OIDC analysis. Teams can now see whether deployment-scoped secrets are stale, unused, or replaceable with OIDC without building a separate GitHub audit script.

StepSecurity

DETECT

Observe SELinux denials and file renames from running workloads

Inspektor Gadget added tracing gadgets for SELinux AVC decisions and file renames, plus IPv6 support in `trace_sni`. That gives responders and platform teams live kernel-level signals for blocked access attempts and filesystem activity without first reproducing the issue under a heavier profiler.

Inspektor Gadget

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

sif

Feedback ⓘ v2026.11.4 NOTES ↗

the blazing-fast pentesting suite.

sif adds scan templates via `-template`, parallel scanning via `-concurrency`, JSON report output via `-json`, config-defined custom detectors, and a TCP module executor.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.11.4 https://github.com/vmfunc/sif.git
# already have the repo? check out this version:
$ git checkout v2026.11.4
```

Load a predefined scan profile with the `-template` flag to run a consistent set of modules without spelling them out each time.

Scan multiple targets in parallel with the `-concurrency` flag.

Write a JSON findings report with `-json` alongside the existing `-sarif` and `-markdown` outputs.

Define custom detectors in config to match responses using `matchers-condition` for OR logic across multiple matchers.

Add a `disable-redirects` per-module HTTP option to prevent individual modules from following redirects.

+7 more

- NEW DETECTION GraphQL introspection detection module
- NEW DETECTION Adminer and phpMyAdmin database panel modules
- NEW DETECTION Grafana, Kibana and Jenkins login panel modules
- NEW DETECTION Active web-vulnerability probe modules
- NEW DETECTION Host-header injection and XXE probe modules
- NEW DETECTION S3, GCS and Azure bucket listing modules
- NEW DETECTION Roundcube and Zimbra webmail fingerprints
- NEW DETECTION PostfixAdmin setup and panel fingerprints
- NEW DETECTION Vaultwarden, Authentik and Synapse version modules
- NEW DETECTION CRLF injection and SSI injection detection modules

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

InGuardians peirates

Feedback ⓘ v1.1.35 NOTES ↗ RANK

Peirates - Kubernetes Penetration Testing tool

> Peirates adds support for authenticating with EKS node tokens, extending its Kubernetes penetration testing capabilities to Amazon EKS clusters.

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

01 **EKS node token authentication** NEW 40

Peirates can now authenticate using EKS node tokens, enabling privilege escalation and pivoting against Amazon EKS clusters.

— Names mechanism but no commands or flags given v1.1.35

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

xalgorix

Feedback ⓘ v4.6.178 NOTES ↗ RANK

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

> xalgorix's latest release adds authentication-bypass detection, JSON request-body contract discovery, and body-based injection for SQL and template injection checks, along with expanded SQL error-based detection coverage.

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

01 Authentication-bypass detection via verify_authbypass NEW

88

The new `verify_authbypass` check re-requests protected routes unauthenticated using equivalent spellings (case fold, percent-encoding, trailing slash, dot and matrix separators) and reports confirmed bypasses under the new `auth-bypass` report class (CWE-287/288/425) at medium or high severity.

— Names check, report class, CWEs, severity but not exact invocation v4.6.178

02 Body injection for SQLi and SSTI checks IMPROVED

83

`verify_sqli` and `verify_ssti` can now inject into JSON or form request bodies in addition to query strings, configured via the new options `inject_in`, `body_type`, and `base_body`.

— Names exact checks and config options, no runnable example v4.6.178

03 JSON body field discovery with probe_body_contract NEW

67

The new `probe_body_contract` tool discovers which JSON body fields a POST/PUT/PATCH endpoint expects by reading the app's own validation errors, classifying each field as required, recognized, rejected as unknown, or type-validated.

— Explains mechanism and classification but no usage example v4.6.178

THINNER COVERAGE BELOW

04 Expanded SQL error signatures for PostgreSQL and SQL Server IMPROVED

35

`verify_sqli` now detects error-based SQL injection on PostgreSQL and SQL Server through expanded engine-specific database error signatures.

— Names engines and check but describes only the outcome v4.6.178

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Rapid7 Metasploit Pro

Feedback ⓘ

November 4, 2026: 5.1.0-2026110301 NOTES ↗ RANK

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

> Metasploit Pro 5.1.0 ships new exploit modules targeting Langflow, an Apache Solr version scanner, and a round of Meterpreter agent improvements including recursive directory creation and a reworked Windows reflective loader.

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

01 Langflow exploit modules for RCE and command injection NEW

65

New exploit module detects and exploits an unauthenticated remote code execution vulnerability in Langflow versions 1.0.0 through 1.9.3, and a second new module detects and exploits an authenticated blind command injection vulnerability in Langflow versions 1.5.0 through 1.9.0.

– Specific version ranges but no module path or command given

November 4, 2026: 5.1.0-2026110301

THINNER COVERAGE BELOW

02 Meterpreter cross-platform agent improvements IMPROVED

52

All Meterpreter agents now support recursive `mkdir -p` capabilities, aligning behaviour across Windows, Python, PHP, and other platforms. Windows Meterpreter gets a major reflective-loader and IAT-import reduction, improved WinAPI and error handling, and MinGW/CMake build refinements, while Python Meterpreter adds improved Windows stat support and FreeBSD process listing.

– Names mechanisms but no flags or commands to invoke November 4, 2026: 5.1.0-2026110301

03 Apache Solr version scanner module NEW

34

New scanner module identifies Apache Solr instances and reports their version.

– Bare description of scanner with no module name or options

November 4, 2026: 5.1.0-2026110301

04 **Pro performance and memory logging** IMPROVED 18

Metasploit Pro adds performance improvements and memory usage logging.

— Generic claim with no mechanism or metric detail November 4, 2026: 5.1.0-2026110301

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ WAF, Bot & API Protection

Arcjet

Feedback ⓘ

Introducing Arcjet's custom AI security models (2026-11-05) BLOG ↗ RANK

Arcjet is the AI agent runtime security platform. Discover the agents running in your organization, enforce policy across every action, prompt, and tool call, and keep the evidence to prove what happened.

> Arcjet overhauls its prompt injection detection with a new in-house model family, deterministic obfuscation decoding, and a specialist two-stage review pipeline for coding agents.

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

01 **Arcjet Thebe 1 and two-stage coding-agent review** NEW

70

Arcjet adds Arcjet Thebe 1, a specialist prompt injection model fine-tuned on coding-agent traffic and available via early access, for teams building coding agents where SQL and shell commands are legitimate but tool output, fetched pages, and installed skills carry injection risk. Coding-agent tool results now go through a two-stage review: a fine-tuned model scans each result and marks suspicious passages, then a decision model reviews those passages together with the agent task, the tool that returned them, and recent session history before issuing a verdict.

— Describes full review pipeline and names early-access entry point.

Introducing Arcjet's custom AI security…

02 **Arcjet Himalia 1 prompt injection model** NEW

65

Arcjet's prompt injection detection now runs on Arcjet Himalia 1, a full fine-tuned general-purpose prompt injection classifier built on Qwen3-0.6B. It returns a binary verdict (benign or injection) rather than a confidence score, so model retrains never silently shift caller behavior and payload-tuning attacks against a confidence threshold are prevented.

— Names model and base architecture but no config or flag to invoke.

Introducing Arcjet's custom AI security…

03 **Deterministic decoders for obfuscated prompt injection** NEW

60

Arcjet now applies deterministic decoders in the request path that recognise `base64`, `ROT13`, `hex`, `leetspeak`, reversed text, fake delimiters, and fake chat-format tokens, and immediately denies those requests before passing decoded plaintext to the model.

— Lists exact encodings handled but no user-facing control surface.

Introducing Arcjet's custom AI security…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

betterleaks

Feedback ⓘ v1.10.0 NOTES ↗

Find leaked secrets everywhere.

betterleaks adds `--unsafe-auto-load-config` and `--trust-local-git-config` flags, changing Git scan isolation defaults.

└──▷ GET THIS VERSION ▾

Homebrew

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

Docker

```
$ docker pull ghcr.io/betterleaks/betterleaks:v1.10.0
```

Go

```
$ go install github.com/betterleaks/betterleaks/v2@v1.10.0
```

git

```
$ git clone --branch v1.10.0
https://github.com/betterleaks/betterleaks.git
# already have the repo? check out this version:
$ git checkout v1.10.0
```

Homebrew lists this version · published under the project's GitHub account · Go module path matches the repo · checked 5 Nov

Wrong command?

Require `--unsafe-auto-load-config` to automatically load `.betterleaks.toml` or `.gitleaks.toml` from the scan target or current directory; explicit `--config`, `BETTERLEAKS_CONFIG`, and `BETTERLEAKS_CONFIG_TOML` are unaffected.

Opt trusted repositories back in to local Git configuration with the new `--trust-local-git-config` flag; local Git scans now isolate repository configuration by default.

NEW DETECTION NuGet API-key detection

UPDATED DETECTION Newer Cloudflare token format support

UPDATED DETECTION Expanded credential validation

└──▷ **BREAKING ON UPGRADE**

! Automatic loading of `.betterleaks.toml` and `.gitleaks.toml` from the scan target or current directory now requires `--unsafe-auto-load-config`; scans that relied on auto-loading will no longer pick up those files without the flag.

! Local Git scans now isolate repository configuration by default; repositories that depended on local Git config being inherited must add `--trust-local-git-config`.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ

CodeQL 2.27.4 improves accuracy and query metadata (2026-11-04) BLOG ↗ RANK

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

> GitHub's CodeQL 2.27.4 release tightens query accuracy across multiple languages and cleans up CLI output, while GitHub Copilot gains a native integration inside Xcode's coding assistant.

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

01 **CodeQL query accuracy and metadata fixes** IMPROVED 64

CodeQL 2.27.4 improves detection and reduces false positives across several queries: `java/log-injection` now recognizes additional sanitizers and sanitizer guards from the JDK, Apache Commons Lang and Text, Guava, OWASP Java Encoder, and Spring Framework; path-injection queries for Java/Kotlin recognize additional substring guards and string replacement/removal operations from Apache Commons Lang and Guava; `rust/uncontrolled-allocation-size`, `rust/log-injection`, and `rust/cleartext-logging` have reduced false positives; `cpp/suspicious-pointer-scaling` no longer flags intentional pointer arithmetic on `std::byte*` buffers; `cpp/unused-local-variable` and `cpp/unused-static-variable` no longer report variables marked `[[maybe_unused]]`; and the CWE-730 tag was removed from regex-related queries across C#, Java/Kotlin, JavaScript/TypeScript, Python, Ruby, and Swift to correct query metadata.

— Names every affected query and guard but no runnable command

CodeQL 2.27.4 improves accuracy and que…

02 **GitHub Copilot in Xcode's coding assistant** NEW 62

GitHub Copilot can now be enabled inside Xcode's built-in coding assistant (public preview) via `Xcode > Settings > Intelligence`, starting with Xcode 27.3 beta on macOS Tahoe 26.6 or later. All Copilot plans are supported.

— Gives exact UI path and version requirements to enable it

GitHub Copilot is now available in Xcod…

THINNER COVERAGE BELOW

03 **CodeQL CLI routes JVM errors to stderr** IMPROVED 36

The CodeQL CLI now writes Java virtual machine warnings and initialization errors to standard error instead of standard output, keeping JSON and other machine-readable output clean.

— Explains the change but gives no flag or command to invoke it

CodeQL 2.27.4 improves accuracy and que…

WAS THIS USEFUL?  

Veracode

Feedback ⓘ Integration updates (2026-11-03) BLOG ↗

Veracode is a cloud-based application security platform that scans code, binaries, and running applications to find and help remediate software vulnerabilities.

Veracode integrations add IaC policy controls, predependency commands, environment diagnostics, scan scheduling, and new filtering views across GitHub, GitLab, Azure DevOps, VS Code, and JetBrains.

Add `predependency_command` to `veracode.yml` in GitHub, GitLab, and Azure DevOps workflow integrations to run commands before Veracode packages a project.

Add `policy` to `veracode.yml` in GitHub, GitLab, and Azure DevOps workflow integrations for IaC scans.

Prevent GitLab merge requests from merging when `break_build_policy_findings` and `break_build_on_error` are both `true`.

Veracode Scan for VS Code adds environment setup checks for connectivity, proxy, DNS, firewall, and TLS issues.

Veracode Scan for VS Code adds a region-change option on the OAuth authentication screen, a Current File filter, file-name grouping for Static Scan Findings, and a Mitigated Findings section in the Flaws That Violate Policy view.

+5 more

NEW SUPPORTED PLATFORM JetBrains IDE versions 2024.2 through 2026.3

NEW SUPPORTED PLATFORM Jira 11 support in Veracode Integration for Jira Server

RETIRED SUPPORTED PLATFORM Jira 9 support end-of-life in Veracode Integration for Jira

NEW INTEGRATION Autopackaging support for Apex, VB6, and Perl in GitHub Workflow Integration

└─▷ **BREAKING ON UPGRADE**

- ! Veracode Azure DevOps Extension v4.0.0 deprecates VeracodeFlawImporter v3 and replaces it with VeracodeFlawImporter v4; existing pipelines must be manually migrated from v3 to v4.
- ! Azure DevOps Server users must upgrade from v3.35.0 to v4.0.0 and manually migrate pipelines.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

cyclonedx-python

Feedback ⓘ v8.0.0 NOTES ↗ RANK

CycloneDX Software Bill of Materials (SBOM) generator for Python projects and environments

> cyclonedx-python 8.0.0 changes how the `environment` command resolves the target interpreter when no explicit path is given, a breaking change for the current major version.

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

01 Auto-detect current interpreter in environment command

75

BREAKING

The `environment` command now probes the current Python interpreter automatically when `-E` or `-S` is passed without a `<python>` argument, e.g. `cyclonedx-py environment -E`. This changes the previous default behaviour, which is why it ships as a major version bump (v8.0.0).

Generate a CycloneDX SBOM from whichever Python interpreter is currently active, without having to supply an explicit interpreter path.

```
$ cyclonedx-py environment -E
```

– Names exact flags and command, but limited mechanism detail beyond behaviour change.

v8.0.0

└──▷ **BREAKING ON UPGRADE**

! The `environment` command now probes the current interpreter when `-E` or `-S` is given with no `<python>` argument, changing the previous default behaviour (major version bump indicates a breaking change in that code path).

WAS THIS USEFUL? 👍 👎

Cycode

Feedback ⓘ

Cycode Extends Agentic Development Security to Citizen Developers (2026-11-04)

BLOG ↗ RANK

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

> Cycode launches Citizen Developer Security, extending its Agentic Development Security Platform to discover unmanaged AI builders, protect their workstations, and scan the code and dependencies their agents generate.

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

01 Citizen developer and unmanaged codebase inventory NEW

55

Cycode now builds a living inventory of every citizen developer and unmanaged codebase, mapping each builder to their team, device, data, and exposure via the Context Intelligence Graph.

— Names the Context Intelligence Graph but no UI path or API

Cycode Extends Agentic Development Secu…

02 AI Guardrails for secrets in prompts and MCP calls NEW

50

AI Guardrails now keep secrets and sensitive data out of prompts, agent file reads, and MCP calls.

— Names three surfaces but no config or enforcement mechanism

Cycode Extends Agentic Development Secu…

03 Scanning of agent-generated code and configuration NEW

50

Cycode now scans agent-generated code, dependencies, secrets, and cloud configuration before it reaches production.

— Lists scan targets but no trigger or integration detail

Cycode Extends Agentic Development Secu…

04 Workstation Protection blocks malicious packages NEW

40

Cycode Workstation Protection now blocks malicious packages before they install on citizen developer devices.

— Clear capability but no detection mechanism detailed

Cycode Extends Agentic Development Secu…

05 **Policy controls for AI tools and MCP servers** NEW 40

Policy controls now govern which AI tools and MCP servers agents can use on a workstation.

— No policy syntax or management UI path given Cycode Extends Agentic Development Secu…

WAS THIS USEFUL? 👍 👎

StepSecurity Feedback ⓘ

Action Secrets: Environment Secrets (2026-11-04) BLOG ↗ RANK

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

> StepSecurity extended the Action Secrets page with full environment secrets coverage (inventory, rotation, usage and OIDC analysis), and shipped a broad round of Harden Runner updates including Red Hat OpenShift support, non-privileged Tetragon mode, reverse-shell host exclusions, automatic GitHub Actions domain allowlisting, and a breaking default change for DNS transaction sweeping.

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

01 OpenShift support for ARC Harden Runner NEW

90

Enable Red Hat OpenShift support in ARC Harden Runner by setting `openshift.enabled=true` in the Helm chart, which creates a SecurityContextConstraints (SCC) object for the harden-runner service account; set `openshift.scc.privileged=false` to use a least-privilege SCC that runs no privileged containers.

Deploy ARC Harden Runner on Red Hat OpenShift without privileged containers, using a least-privilege SCC.

yaml

```yaml
openshift:
  enabled: true
  scc:
    privileged: false
```

— Exact Helm keys with a runnable config example.

Harden Runner Releases: November 4, 202…

02 Environment secrets support in Action Secrets NEW

83

The secrets inventory API and MCP server now return environment secrets alongside repository and organization secrets, gated behind a new Environments: Read permission for the StepSecurity GitHub App (a new API endpoint reports whether this permission is granted; existing installations need an organization owner to approve it). The Action Secrets page adds an Environment Secrets tab showing every environment secret by

repository and environment — filterable by environment — with rotation age, usage status, referencing workflows, OIDC replaceability, and a lock icon marking protected environments with their protection rules. Environment secrets are now covered by the `Secrets should be rotated periodically` control using the same 180-day threshold; a new `Possible` usage status flags secrets used only by jobs that choose their environment dynamically at run time (for example `environment: ${{ inputs.env }}`) so they aren't deleted or rotated without review; usage tracking now follows GitHub's resolution order (environment, then repository, then organization), marking overridden repository secrets `Unused` with an `Overridden in` note; and OIDC replaceability recommendations are now computed per secret using only the workflow steps that actually receive it.

.
.
.

— Rich mechanism and named permission/statuses, but no runnable command.

`Action Secrets: Environment Secrets (20…`

---

03 **Host exclusions for reverse shell detection** NEW 83

Exclude specific hosts from Reverse Shell detection using the new opt-in `env.lockdownExcludeHostRegex` list, so `/dev/tcp/<host>/<port>` connections to matched hosts (such as in-cluster service probes) no longer trigger alerts.

Suppress reverse shell alerts for in-cluster health probes that use `/dev`/tcp connections to internal service hosts.

```yaml
env:
  lockdownExcludeHostRegex:
    - '^healthcheck\.internal\.svc\.cluster\.local$'
    - '^10\.96\..*'
```

— Named flag and pattern, with a concrete suppression example.

`Harden Runner Releases: November 4, 202…`

---

04 **Custom volume mounts for Harden Runner pods** NEW 75

Mount additional volumes — such as a corporate CA certificate for TLS-inspecting proxies like Zscaler — using the new Helm values `env.volumes` and `env.volumeMounts`.

Mount a corporate CA certificate into the runner for environments with a TLS-inspecting proxy such as Zscaler.

```yaml
env:
  volumes:
    - name: corporate-ca
      secret:
        secretName: corporate-ca-cert
  volumeMounts:
    - name: corporate-ca
      mountPath: /etc/ssl/certs/corporate-ca.crt
      subPath: ca.crt
```

— Named config keys plus a runnable YAML example.

Harden Runner Releases: November 4, 202…

05 **Non-privileged Tetragon BPF mount mode** NEW 70

Run Tetragon without privileged mode by setting the new Helm value `tetragon.bpfMountPropagation` to `HostToContainer` (the default remains `Bidirectional`).

— Exact config key and values but no example snippet.

Harden Runner Releases: November 4, 202…

06 **Block policy confirmation in Cilium-only mode** IMPROVED 70

ARC Harden Runner now waits in Cilium-only mode until the Cilium agent reports a block policy as active on every tracked pod before marking it applied, falling back to previous behavior on timeout or missing socket; the Helm chart mounts the Cilium agent socket when `dnsMode` uses Cilium.

— Describes mechanism and config condition but no exact flag.

Harden Runner Releases: November 4, 202…

THINNER COVERAGE BELOW

07 **DNS transaction sweep enabled by default** BREAKING 55

`env.enableDnsTransactionSweep` now defaults to `true`; any setup relying on the previous default of `false` will change behavior on upgrade.

— Exact config key and clear migration note for upgraders.

Harden Runner Releases: November 4, 202…

08 **Automatic GitHub Actions domain allowlisting** IMPROVED 45

GitHub Actions domains from GitHub's meta API are now always allowed automatically by block policies and refresh every hour, removing the need to manually allowlist them.

— Explains behavior but requires no reader action.

Harden Runner Releases: November 4, 202…

---

09 **HTTP/2 traffic monitoring in Self-Hosted VM Harden Runner**

IMPROVED 30

Self-Hosted VM Harden Runner v1.9.6 improves HTTPS monitoring for HTTP/2 traffic.

— Named version but mechanism is unspecified. Harden Runner Releases: November 4, 202…

└──▷ **BREAKING ON UPGRADE**

! `env.enableDnsTransactionSweep` now defaults to `true`; any setup relying on the previous default of `false` will change behavior on upgrade.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-11-02 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds a Malcat HTTP API bridge, sharded file storage with deduplication, HTTP/3 support in Nginx, and Proxmox API token authentication

└──▷ TRY IT

Start the Malcat bridge so CAPE can reach Malcat features without sharing a venv — first generate a key, store the license, then serve on localhost.

```
$ python malcat_bridge.py genkey
python malcat_bridge.py setlicense <your-malcat-license-key>
python malcat_bridge.py serve --host 127.0.0.1 --port 8090
```

Serve the Malcat bridge with TLS so the channel is encrypted without putting nginx in front immediately.

```
$ python malcat_bridge.py gencert
python malcat_bridge.py serve --host 127.0.0.1 --port 8090 --tls-cert
cert.pem --tls-key key.pem
```

Allow Proxmox API token authentication and configurable TLS verification for Proxmox machinery connections.

CAPEv2 adds a `malcat_bridge.py` HTTP API bridge that exposes Malcat features to CAPE over the network, with hashed API key authentication via `X-API-Key` header, optional HTTPS via `--tls-cert` and `--tls-key`, and subcommands `genkey`, `setlicense`, `gencert`, and `serve`.

CAPEv2 adds sharded file directories, deduplication, and NFS sync to storage via the new storage subsystem.

CAPEv2's Nginx installer now enables HTTP/3 (QUIC) configuration.

CAPEv2's web UI now shows an immediate static analysis preview on the submission status page.

+3 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Sublime Security

Feedback ⓘ

Platform v1.90.180 (Prompt Injection Detection) (2026-11-04) BLOG ↗

Sublime Security is an email security platform that uses detection-as-code rules to identify and block phishing, business email compromise, and other malicious messages.

Sublime Security adds a Prompt Injection card in Attack Overview and a Prompt Injection tag on the message details page to surface ASA-confirmed attempts.

NEW DETECTION A Prompt Injection card in Attack Overview now opens the Message List View filtered to ASA (Autonomous Security Analyst) confirmed prompt injection attempts.

The message details page now shows a Prompt Injection tag next to the ASA verdict for confirmed attempts.

WAS THIS USEFUL? 👍 👎

Sophos Email Security

Feedback ⓘ

2026.44Released November 2, 2026 NOTES ↗ RANK

Sophos Email Security is a cloud-based service that filters inbound and outbound email to block spam, phishing, malware, and data loss.

> Sophos Email Security adds Image Analysis, letting admins detect unwanted images in email and act on them through policy.

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

01 **Image Analysis for unwanted images** NEW 74

Sophos Email Security now supports Image Analysis to detect unwanted images in email messages and take action based on their category. It is configurable in both Email Security and Data Control policies, with results visible in reports and dashboards. To enable it, go to Email Security > Policies > Email Security Policy, select a policy, navigate to the Image Analysis section, enable detection, choose image categories to act on, and set an action such as quarantine or block; results appear under Reports > Email Security dashboards.

Enable Image Analysis on inbound mail to block or quarantine messages containing unwanted image categories.

📍 In Sophos Central, go to Email Security › Policies › Email Security Policy, select your policy, navigate to the Image Analysis section, enable detection, choose the image categories to act on, and set the desired action (e.g., quarantine or block). Repeat for Data Control policies if needed. Review results under Reports › Email Security dashboards.

— Names exact UI path and policy areas, but no categories/limits listed

2026.44Released November 2, 2026

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Vectra AI

Feedback ⓘ Rapid Release Improvements (2026-11-04)

BLOG ↗

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

NEW DETECTION EDR Communication Disruption Detection

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ Secrets and end-of-life checks for Azure and GCP:

NOTES ↗ RANK

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

> Plerion added automated secret detection for GCP Cloud Functions and Compute Engine startup scripts, plus a new check that flags Azure Functions running end-of-life runtimes.

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

01 **GCP Cloud Function secret detection** NEW 55

Plerion now raises findings via `PLERION-GCP-18` when a Cloud Function runtime or build environment variable holds something that looks like an API key, token, or private key — previously this was a manual recommendation only.

— Names the check id and before/after but no remediation steps.

Secrets and end-of-life checks for Azur…

02 **Azure Functions end-of-life runtime detection** NEW 50

Plerion now flags Azure Functions on end-of-life runtimes via `PLERION-AZURE-463`, covering .NET, Java, Node.js, PowerShell, and Python.

— Names check id and languages but no mechanism or remediation detail.

Secrets and end-of-life checks for Azur…

WAS THIS USEFUL? 👍 👎

Upwind

Feedback ⓘ

Upwind + ServiceNow VR: Only real findings reach the ticket queue (2026-11-04)

BLOG ↗ RANK

Upwind is a runtime-powered CNAPP that protects cloud workloads by combining posture, detection and response, and identity security with real-time runtime context.

> Upwind launched a ServiceNow Vulnerability Response integration that files findings as native Vulnerable Items instead of generic ITSM tickets, with automated and one-off ticketing paths.

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

01 ServiceNow Vulnerability Response integration NEW

91

Connect Upwind to ServiceNow VR via Components → Integrations → ServiceNow VR by entering your instance URL and OAuth client credentials and clicking Provision, which auto-creates the staging table and related objects in ServiceNow with nothing to install on that side. Findings are sent as native Vulnerable Items (VIT) linked to their CVE record from the NVD library, matched to the Configuration Item, and tagged `Source = Upwind`. The When a new vulnerability is found workflow trigger paired with the Create a ServiceNow Vulnerable Item action files a VIT automatically for every matching finding, filterable by severity, EPSS, risk category, or any selector; individual findings can also be sent one-off from the finding page, after which the VIT number, state, assignment group, and assignee appear on the finding linking back to ServiceNow.

Set up the ServiceNow VR integration for the first time so Upwind can start filing Vulnerable Items in your remediation team's queue.

📍 In Upwind, go to Components → Integrations → ServiceNow VR. Enter your ServiceNow instance URL and OAuth client credentials, review the field mapping, then click Provision. Upwind will create its staging table and related objects in ServiceNow — nothing needs to be imported or installed on the ServiceNow side.

Automatically create a Vulnerable Item for every new high-severity finding so remediation queues stay current without manual ticketing.

📍 In Upwind Workflows, create a new workflow with the trigger When a new vulnerability is found, add a filter for severity (e.g. Critical or High) or EPSS score, then set the action to Create a ServiceNow Vulnerable Item. Save and enable the workflow.

— Names exact UI path, flags, trigger/action pair, and field mappings.

Upwind + ServiceNow VR: Only real findi…

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Inspektor Gadget

Feedback ⓘ v0.58.0 NOTES ↗ CODE ↗

RANK

Inspektor Gadget is a set of tools and framework for data collection and system inspection on Kubernetes clusters and Linux hosts using eBPF

> Inspektor Gadget v0.58.0 adds two new tracing gadgets for SELinux AVC decisions and file renames, extends `trace_sni` to IPv6, and improves OpenTelemetry metrics and profiling output with a snapshot mode and a label-sanitizing flag, alongside a breaking change to `trace_fsslower` JSON fields.

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

01 New trace_selinux gadget for SELinux AVC decisions NEW

85

New `trace_selinux` gadget traces SELinux AVC decisions, including security contexts, decoded permission names from the running kernel, and optional kernel and user stacks.

Trace SELinux AVC denials on a Linux host to identify policy violations with security contexts and decoded permissions.

```
$ sudo ig run trace_selinux:latest
```

— Names mechanism and includes a runnable command v0.58.0

02 OTel profiling label sanitization flag NEW

83

New `--otel-profiles-sanitize-labels` flag sanitizes OpenTelemetry profiling sample attribute names for dashboard compatibility by replacing dots with underscores.

Run the eBPF profiler with sanitized label names so sample attributes are compatible with existing OTel dashboards that reject dots in label names.

```
$ sudo ig run --otel-profiles-sanitize-labels <gadget>:latest
```

— Named flag with exact behaviour and command v0.58.0

03 Breaking change to trace_fsslower JSON fields BREAKING

80

In `trace_fsslower`, the `delta_us` field in JSON output is replaced by numeric `delta_raw` (in nanoseconds) and formatted `delta`; the `size` field becomes a formatted string with the numeric value available as `size_raw`.

— Exact field-level migration detail for JSON output v0.58.0

---

04 **New `trace_rename` gadget for file renames** NEW 75

New `trace_rename` gadget reports successful file renames with optional absolute paths and rename flags.

Trace file rename operations on a Linux host to audit which files are being renamed and by which processes.

```
$ sudo ig run trace_rename:latest
```

— Clear description plus runnable command v0.58.0

THINNER COVERAGE BELOW

05 **IPv6 support in `trace_sni`** IMPROVED 50

`trace_sni` now extracts TLS server names from IPv6 Client Hello packets in addition to IPv4.

— Describes change but no example or command given v0.58.0

---

06 **OpenTelemetry metrics snapshot mode** NEW 50

OpenTelemetry metrics now support an opt-in snapshot mode that removes stale series for processes or entities that no longer exist while preserving updates for those still active.

— Explains behaviour but lacks concrete config or command v0.58.0

└──▷ **BREAKING ON UPGRADE**

! In `trace_fsslower`, the `delta_us` field in JSON output is replaced by numeric `delta_raw` (in nanoseconds) and formatted `delta`; the `size` field becomes a formatted string with the numeric value available as `size_raw`.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ Secure SaaS - November 4, 2026 BLOG ↗

RANK

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

> Sysdig Secure's CLI Scanner v1.32.0 broadens standalone binary scanning and OS platform coverage, sharpens vulnerability detection logic, and introduces two breaking changes around Fedora package matching and older on-premises collectors.

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

01 **Standalone binaries excluded from SBOM on older on-prem collectors**

85

BREAKING

When scanning images against on-premises collectors 7.5.0 to 7.9.x that contain standalone binary applications, the CLI Scanner now excludes applications from the SBOM and scan results. Upgrade the on-premises backend to 7.10.0 or later to restore application inclusion.

— Names exact affected versions and clear upgrade path Secure SaaS - November 4, 2026

02 **PHP version detection via PHP_VERSION environment variable**

IMPROVED 60

PHP detection in official Docker Hub `php` images now also reads the `PHP_VERSION` environment variable.

— Names exact env var reader can verify in images Secure SaaS - November 4, 2026

THINNER COVERAGE BELOW

03 **New OS platform support for scanning** NEW 55

Image scanning now supports SUSE Linux Enterprise Server 16, openSUSE Leap 16, and Ubuntu 26.10, including end-of-life information for SUSE Linux Enterprise Server 16.

— Names exact OS versions, no usage steps Secure SaaS - November 4, 2026

04 **Standalone binary scanning for Python, PHP, Chrome Headless Shell, Muse Code**

NEW 53

CLI Scanner v1.32.0 adds standalone binary scanning support for Python, PHP, Chrome Headless Shell, and Muse Code applications.

— Names targets but no mechanism or command given Secure SaaS - November 4, 2026

05 **Fedora vulnerability matching changed to source package**

50

BREAKING

After upgrading, scans of Fedora images may show additional vulnerabilities because Fedora subpackages are now matched against their source package instead of the previous matching logic.

— Explains before/after but no remediation steps Secure SaaS - November 4, 2026

06 **Claude Code detection during image scanning** NEW

40

The scanner can now detect Claude Code installed via its native installer during image scanning.

— States capability but no further detail on scope Secure SaaS - November 4, 2026

07 **Exploitable vulnerabilities marked without publication date**

IMPROVED

40

Vulnerabilities with a known exploit are now marked exploitable even when the exploit has no publication date or reference link.

— Describes logic change but no configuration or trigger Secure SaaS - November 4, 2026

└──▷ ALSO FROM THESE RELEASES

.
.
.

└──▷ BREAKING ON UPGRADE

! After upgrading, scans of Fedora images may show additional vulnerabilities because Fedora subpackages are now matched against their source package instead of the previous matching logic.

! When scanning images against on-premises collectors 7.5.0 to 7.9.x that contain standalone binary applications, the CLI Scanner now excludes applications from the SBOM and scan results. Upgrade the on-premises backend to 7.10.0 or later to restore application inclusion.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ

Find misconfigured SIEM log sources from your own tools with the API NOTES ↗

RANK

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

> Huntress shipped an API endpoint for querying misconfigured SIEM log sources and added a self-service Troubleshooting tab to PSA integration configuration pages.

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

01 **Troubleshooting tab for PSA integrations** NEW 95

A new Troubleshooting tab on PSA integration Configure pages lets account admins diagnose common configuration issues without filing a support ticket, available for ConnectWise, Autotask, Syncro, HaloPSA, and Kaseya BMS. It includes a Permission Test that runs query checks against the PSA to confirm API credentials have required access, and for ConnectWise, Autotask, and HaloPSA a Billing Sync Dry Run that executes the full sync logic without writing changes so results can be previewed before committing.

Verify your ConnectWise API credentials have the right permissions before going live, without filing a support ticket.

> 📍 In the Huntress portal, go to Account Settings › PSA Integrations › ConnectWise › Configure › Troubleshooting, then run the Permission Test query to confirm your API credentials are correctly scoped.

Preview a full billing sync against HaloPSA to catch setup errors before any data is written.

> 📍 In the Huntress portal, go to Account Settings › PSA Integrations › HaloPSA › Configure › Troubleshooting, then run Billing Sync Dry Run to see the full sync output without committing changes.

— Names exact UI path, PSAs, and the two diagnostic tools with their behavior

New Troubleshooting tab for PSA integra…

02 **API access to misconfigured SIEM log sources** NEW 70

A new API endpoint lets you query misconfigured SIEM log sources, returning agent ID, hostname, organization, status (`unresolved` or `resolving`), and the full required audit policy (category, name, success/failure) — the same data shown on the

Misconfigured Sources tab in the portal. Results can be filtered by `organization_id` and paginated to integrate the data into your own tooling.

— Fields and filter named but no explicit endpoint path or method given

Find misconfigured SIEM log sources fro…

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

New Images view in Container Security inventory (2026-10-30) 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 introduces a dedicated Images view in Container Security Inventory, letting teams browse, inspect, and compare container images from both runtime and pre-runtime scans.

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

01 **Images view in Container Security Inventory** NEW 77

A new Images section appears under Container Security Inventory, navigated via Cloud Security → Container Security → Inventory → Images → All images or a specific registry, listing each container image with vulnerability counts and running workloads. Selecting an image opens a detail page covering overview, workloads, vulnerabilities, malware, secrets, packages, and scan history, and the same view surfaces images discovered pre-runtime via TrendAI Artifact Scanner (TMAS) alongside runtime-discovered images. Two images can be selected for a side-by-side comparison showing introduced, fixed, and unchanged vulnerabilities, malware, and secrets.

Compare two container images in the inventory to identify newly introduced vulnerabilities, malware, or secrets before promoting a new image to production.

> 📍 In the console, go to Cloud Security → Container Security → Inventory → Images → All images. Select the checkbox next to two images, then review the side-by-side comparison showing introduced, fixed, and unchanged vulnerabilities, malware, and secrets.

— Exact navigation path and comparison workflow given, but no limits or numbers.

New Images view in Container Security i…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

World Monitor

Feedback ⓘ v2.11.8 NOTES ↗ RANK

Real-time global intelligence dashboard. AI-powered news aggregation, geopolitical monitoring, and infrastructure tracking in a unified situational awareness interface

> World Monitor's v2.11.8 release adds Pro license activation for self-hosted dashboards and a new daily Thailand farm-price data feed.

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

01 **Thailand farm-price data from NABC dashboard** NEW 40

World Monitor now reads Thailand's daily farm prices from the NABC dashboard as a new data source.

— Names the data source but no mechanism or UI path given. v2.11.8

02 **Pro license activation on self-hosted dashboards** NEW 25

World Monitor adds Pro license activation for self-hosted dashboards.

— Only mentioned in summary, no detail on how activation works. v2.11.8

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Splunk Security Content

Feedback ⓘ v6.10.0 NOTES ↗

RANK

Splunk Security Content

> Splunk Security Content v6.10.0 ships 14 new detections and a new Node2Shell analytic story spanning Node.js, PHP, Linux auditd, Azure AD and Windows, while retiring ten legacy detections previously flagged for deprecation.

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

01 **Removal of ten deprecated detections** BREAKING 74

Ten detections are removed in v6.10.0, previously announced for deprecation in v6.8.0: Child Processes of Spoolsv exe, Detect HTML Help Renamed, Detect Renamed 7-Zip, Detect Renamed PSExec, Detect Renamed RClone, Detect Renamed WinRAR, Detect WMI Event Subscription Persistence, Detect mshta renamed, MacOS - Re-opened Applications, and Suspicious MSBuild Rename. Users should review the deprecated analytics documentation and migrate to the listed replacement detections before upgrading.

— All ten removed detections named with a migration pointer. v6.10.0

02 **PHP worker segfault linked to multipart POST** NEW 70

A new analytic, PHP Worker Segmentation Fault After Multipart POST, correlates Apache segmentation-fault errors with multipart POST requests handled by the same worker process to link crafted-request crashes to the triggering requests.

— Named detection with clear correlation mechanism described. v6.10.0

03 **Three new Windows persistence and evasion detections** NEW 64

Three new Windows detections cover scheduled task XML built via console echo commands (Windows Scheduled Task XML Built Via Console Echo), user-profile binaries terminating processes via taskkill (Windows User Profile Binary Terminating Process Via Taskkill), and Chromium sessions restored with the crash bubble suppressed (Windows Chromium Session Restored With Crash Bubble Suppressed).

— All three detections named verbatim with their triggers. v6.10.0

THINNER COVERAGE BELOW

04 **Five new Linux auditd detections** NEW 56

Five new Linux auditd analytics cover inline PHP code execution, process image change under the same PID, raw socket creation from user-writable paths (Linux Auditd SOCK RAW EAFNOSUPPORT From User Writable Path), service crash and restart bursts, and a web or application server spawning a shell.

— One exact detection name given; others described only in prose. v6.10.0

05 **Compiled C source staged to tmp detection** NEW 54

A new detection, Linux Compiled C Source To Tmp , flags compiled C source written to temporary directories, improving visibility into post-exploitation staging on Linux hosts.

— Named detection with purpose stated but no further mechanism. v6.10.0

06 **Node2Shell analytic story for Node.js Brotli exploitation**

NEW 53

A new analytic story covers Node2Shell, an out-of-bounds write exploitation technique in the Node.js Brotli decoder, with endpoint and web signals for exploitation attempts and resulting command execution.

— Names the vulnerability and detection scope but no exact detection names. v6.10.0

07 **Fidelity updates to nine existing analytics** IMPROVED

48

Updates nine existing analytics to improve fidelity and broaden coverage: HTTP C2 user agents, large HTTP POST requests that trigger server errors, Linux ngrok reverse proxy usage, React and Next.js child processes, PHP memory exhaustion, serialized PHP upload-filename detection, web server shell-spawn, Windows non-shell command-line tool execution, and Windows Remote Management spawned processes.

— Lists all nine areas touched but no specifics of what changed. v6.10.0

08 **Four new Azure AD identity-abuse analytics** NEW 46

Four new Azure AD analytics detect federated identity credentials added to service principals, guest users changed to members, suspicious device registrations, and Temporary Access Pass creation.

— Lists all four behaviors but no exact detection titles given. v6.10.0

└──▷ **BREAKING ON UPGRADE**

! Ten detections are removed in v6.10.0 (announced in v6.8.0): Child Processes of Spoolsv exe, Detect HTML Help Renamed, Detect Renamed 7-Zip, Detect Renamed PSExec, Detect Renamed RClone, Detect Renamed WinRAR, Detect WMI Event Subscription Persistence, Detect mshta renamed, MacOS - Re-opened Applications, and Suspicious MSBuild Rename. Review the deprecated analytics documentation and migrate to the listed replacement detections before upgrading.

WAS THIS USEFUL? 👍 👎

Tenzir Feedback ⓘ

Tenzir Node v6.21.0: Faster DuckDB, ClickHouse, and Splunk queries (2026-11-04)

BLOG ↗ RANK

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

> Tenzir Node v6.21.0 speeds up federated queries by pushing filtering, sorting, and limiting work down into the underlying databases and by passing typed results back from Splunk instead of raw strings.

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

01 Filter, sort, and limit pushdown for SQL connectors IMPROVED

70

Tenzir now pushes filters, sorting, and limits after summaries down into `DuckDB`, `ClickHouse`, `MySQL`, and Microsoft SQL Server so the database itself returns only the rows the pipeline keeps. Filters also now pass through `ocsf_cast` and field-selection functions instead of blocking on them, and operators placed after `from_duckdb` now use native DuckDB semantics instead of generic ones.

— Names four engines and two operators but no new flags to invoke

Tenzir Node v6.21.0: Faster DuckDB, Cli…

THINNER COVERAGE BELOW

02 Typed result pass-through for Splunk searches IMPROVED

40

Splunk searches now run faster because results are passed back through the pipeline as typed values instead of raw strings.

— Explains mechanism briefly but offers no configurable surface

Tenzir Node v6.21.0: Faster DuckDB, Cli…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

AI Function tasks for jobs (Beta) (2026-11-04) BLOG ↗ RANK

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

> Databricks introduced a Beta job task type that runs AI Functions like `ai_parse_document`, `ai_extract`, and `ai_classify` directly on jobs without writing SQL.

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

01 **AI Function tasks for jobs (Beta)** NEW 80

A new job task type lets you run `ai_parse_document`, `ai_extract`, or `ai_classify` on files in a Unity Catalog volume or rows of a table without writing SQL. The task supports incremental processing of new or changed inputs and writes results to a Unity Catalog table for downstream tasks. In the Jobs UI, add a task, select AI Function as the task type, choose the function (e.g. `ai_classify`), point it at an input table or volume, and specify a Unity Catalog output table.

Add an AI Function task to a job to classify rows in a table without writing SQL, writing results to a Unity Catalog table for downstream tasks.

> 📍 In the Jobs UI, open your job and click Add task, then select AI Function as the task type, choose `ai_classify` as the function, point it at your input table, and specify a Unity Catalog output table.

— Names exact functions and UI steps but feature still Beta

AI Function tasks for jobs (Beta) (2026…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ updated · Nov 5, 2026 NOTES ↗

RANK

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

> NodeZero's Remote Access Tool now supports AARCH64 architecture, extending EDR protection testing to ARM-based hosts.

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

01 **AARCH64 support for Remote Access Tool** NEW 45

NodeZero's Remote Access Tool (RAT) now runs on AARCH64 architecture, enabling EDR protection testing on ARM-based hosts.

— Names the architecture and tool but no setup detail. snapshot-20261105

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Qualys

Feedback ⓘ changelog · Nov 4, 2026 BLOG ↗ RANK

Qualys is a cybersecurity blog covering software, services, and risk insights. Stay updated with the latest on threat detection, compliance, and vulnerability management.

> Qualys shipped a containerized Cloud Agent for Kubernetes, extending host Vulnerability Management and Policy Audit to immutable OS nodes and pairing with TotalCloud Kubernetes and Container Security for full-stack coverage.

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

01 **Containerized Cloud Agent deployed as Kubernetes DaemonSet**

NEW 84

Run Qualys Cloud Agent on every Kubernetes node as a DaemonSet by pulling `qualys/cloud-agent-container-x64` (Intel/x64) or `qualys/cloud-agent-container-arm` (ARM) from Docker Hub, with no host package manager required. The agent can be deployed and upgraded via a Qualys Helm chart published on Artifact Hub alongside deployment YAML manifests to fit existing GitOps workflows. A single container image spans RPM-based, Debian-based, and container-optimized Linux distributions across Intel/x64 and ARM, eliminating per-OS image variants, and FIPS-compliant images are available for environments requiring Federal Information Processing Standards compliance.

— Names exact image tags, Helm chart and Artifact Hub, but no CLI command shown

2026-11-04-2d658fc3

02 **Vulnerability Management and Policy Audit on immutable OS nodes**

NEW 63

Host Vulnerability Management and Policy Audit now cover container-optimized immutable operating systems — AWS Bottlerocket, Google Container-Optimized OS, and Red Hat CoreOS — without modifying the host image. Policy Audit adds CIS-certified policies for AWS Bottlerocket and Google Container-Optimized OS nodes, with findings prioritized by Qualys TruRisk alongside the rest of the asset estate.

— Names specific OS targets and CIS/TruRisk but no config or steps 2026-11-04-2d658fc3

THINNER COVERAGE BELOW

03 **Full-stack Kubernetes coverage via TotalCloud KCS pairing**

NEW 53

Pairing the containerized agent with Qualys TotalCloud Kubernetes and Container Security (KCS) delivers full-stack coverage across the host OS, running containers, and the Kubernetes control plane — including drift detection, runtime security via eBPF sensors, and Kubernetes Security Posture Management (KSPM).

.

— Describes integration scope and mechanisms but no setup instructions

2026-11-04-2d658fc3

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

syzkaller

Feedback ⓘ changes since 2026-11-03 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds riscv64 and arm64 compat target support in syz-agent and AI job creation for riscv64 bugs

syz-agent now supports riscv64 and arm64 compat targets, including arm-on-arm64 compat target selection via `TargetVMArch` in target config.

The dashboard now creates AI jobs for riscv64 bugs, expanding automated AI workflow coverage beyond previously supported architectures.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ

Delegate certificate management to Google Play apps on Android NOTES ↗ RANK

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

> Iru now lets admins delegate certificate installation and selection on managed Android devices to a trusted Google Play app.

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

01 Certificate management delegation for Google Play apps on Android

NEW 75

A new Certificate management section on the Google Play app Library Item lets admins delegate certificate installation, management, and certificate selection to a Google Play app on managed Android devices. In the admin console, under Library > Google Play app > [your PKI app] > Certificate management, admins can enable Install and manage certificates and/or Select certificates for other apps for a third-party PKI app.



Allow a third-party PKI app to install and select certificates for other apps on managed Android devices.

> 📍 In the admin console, go to Library › Google Play app › [your PKI app] › Certificate management, then enable Install and manage certificates and/or Select certificates for other apps, and save.

— Names exact UI path and toggle options; lacks broader mechanism details

Delegate certificate management to Goog…

NEW SUPPORTED PLATFORM App preservation during Return to Service for iOS/iPadOS 26

> Preserve managed apps with Return to Service NOTES ↗ LATEST STABLE

> Iru adds an option to preserve managed apps across Return to Service erases, so redeployed devices are ready for the next user without re-downloading apps.
>
> Enable app preservation through a Return to Service erase by turning on 'Allow app preservation during Return to Service' in the Automated Device Enrollment Library Item for iPhone and iPad, available on iOS and iPadOS 26 or later.

WAS THIS USEFUL? 👍 👎

NinjaOne

Feedback ⓘ announced · Nov 3, 2026 BLOG ↗

Ninjaone is a improve IT efficiency with an autonomous patch management solution. Automate updates and secure endpoints across your organization with NinjaOne today.

NinjaOne adds Chat, a real-time messaging channel for technicians to message users on managed Windows devices from the web app, integrated with Ticketing.

NEW SUPPORTED PLATFORM Start a chat with the user logged in to any managed Windows device directly from the device page or device search in the NinjaOne web app, without opening a remote session or creating a ticket.

Run multiple simultaneous chat sessions and switch between them in the Chat drawer, with typing indicators and read receipts showing conversation status.

Hand off an active chat session to a second technician when one joins the conversation.

Download chat transcripts from Activities once a session ends, with all conversations automatically captured in the associated ticket via the Ticketing integration.

NEW SUPPORTED PLATFORM ChromeOS and ChromeOS Flex device support

▾ **1 more release in this issue** · 2026-11-02 → 2026-11-03

announced · Nov 2, 2026 BLOG ↗

NinjaOne 16.0 adds ChromeOS support: unified inventory, multi-tenant reporting, device actions, and `/v2/devices` API access for Chromebook fleets.

ChromeOS and ChromeOS Flex devices now appear in the existing `/v2/devices` API routes, so MSPs and IT teams can pull Chromebook data the same way they pull Windows, Mac, and mobile devices.

Extend the existing Google Workspace integration to sync ChromeOS and ChromeOS Flex devices into the shared NinjaOne inventory alongside Windows, Mac, and mobile endpoints, with online and offline status per device.

Connect multiple Google Workspace tenants and report across all of them from a single NinjaOne view, covering multi-tenant MSP and district environments without logging into each Workspace console.

Control device onboarding with auto-approve or manual approval, and merge discovered devices into existing records by serial number.

Run device actions — reboot, set volume, erase user data, power wash a Chromebook, or open a Chrome Remote Desktop session — directly from NinjaOne without switching consoles.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-11-05) BLOG ↗

RANK

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

> JumpCloud's latest release adds compliance-grade Access Certification exports and continuous SIEM ingestion of identity events, new Agent Uninstall Protection for managed devices, MCP-based app and workflow management through AI clients, and new per-application App Catalog and Windows policy controls.

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

01 Agent Uninstall Protection for managed devices NEW

95

Agent Uninstall Protection is on by default for devices running JumpCloud agent v2.179.2 or later, requiring a per-device maintenance token to uninstall the agent on macOS, Windows, and Linux. The Disable Agent Uninstall Protection policy can turn off the token requirement on specific devices for maintenance, troubleshooting, or imaging; each device's status (On , Off , or Failed) is visible from device details or the Devices list, and the maintenance token, which rotates 60 minutes after viewing, is retrievable from device details by a Manager role or higher.

Retrieve a device maintenance token to authorize a one-time agent uninstall on a protected device.

> 📍 In the Admin Portal, go to Device Management › Devices, open the device details, and select Agent Uninstall Protection to view the maintenance token (rotates 60 minutes after viewing).

— Exact policy name, roles, rotation timing and UI path Release Notes 2026 (2026-11-05)

02 App Catalog update notification and self-service controls

IMPROVED

85

The App Catalog now lets admins control update notifications per application via Disable end user update notifications when adding or editing an app, while auto-updates still force after seven days when Auto-Update is enabled; users can also start available updates for Required applications directly from the Updates section in the Self Service App Catalog.

— Exact config flag, timing, and UI section named Release Notes 2026 (2026-11-05)

03 **OAuth authorization for workflow connectors** NEW 80

Workflow connectors can be authorized using OAuth 2.0 Client Credentials or Authorization Code with PKCE directly from the Admin Portal, with credentials stored in Secrets Manager and access tokens refreshed automatically.

— Names OAuth flows and storage mechanism, UI path only Release Notes 2026 (2026-11-05)

04 **Browser discovery in AI & SaaS Management** NEW 80

Browser discovery in AI & SaaS Management can be started from a JumpCloud-managed device or by deploying the JumpCloud Go extension through a third-party MDM, using a Client ID and Client secret copied from Settings.

Deploy JumpCloud Go extension through a third-party MDM to enable browser-based SaaS discovery without JumpCloud device management.

> 📍 In the Admin Portal, go to Settings, copy the Client ID and Client secret, then add them to your third-party MDM's JumpCloud Go extension distribution policy and deploy.

— Exact step-by-step deployment path with named credentials

Release Notes 2026 (2026-11-05)

05 **SIEM integration for identity events** NEW 75

JumpCloud identity events can be ingested into SIEM platforms continuously via API key, in RAW or OCSF-normalized formats, or pulled via asynchronous historical export jobs covering up to 7-day windows across a 13-month retention lookback from Directory Insights.

— Concrete formats and retention numbers, but no endpoint given

Release Notes 2026 (2026-11-05)

06 **MCP-only application setup** NEW 75

Adding an application now supports capability-based setup, letting admins select SSO, SCIM, and/or MCP individually so MCP-only apps can bypass SSO configuration and be made available through the AI Gateway. MCP-only apps can also be created directly via Manage MCP Servers and assigned to user groups for access control without configuring SSO or creating a bookmark.

— Named UI path and capability toggles but no exact command

Release Notes 2026 (2026-11-05)

07 **User-level Windows registry and MDM policies** NEW 75

Two new user-level Windows policies ship: the User Level Windows Custom Registry Keys Policy, which manages specific registry values under `HKEY_CURRENT_USER` at sign-in without MDM enrollment or a device restart, and the User Level Custom Windows

MDM (OMA URI) Policy, which configures custom MDM settings assignable to users or user groups on MDM-enrolled Windows devices.

— Named policies and registry path but no step-by-step Release Notes 2026 (2026-11-05)

08 **Access Certification export with tamper-proof hashing** NEW

60

Access Certification campaign snapshots can be exported in CSV, JSON, or PDF formats, each secured with SHA-256 document hashing to support compliance audits.

— Names formats and hashing but no exact export path Release Notes 2026 (2026-11-05)

09 **Workflow creation via MCP-enabled AI client** NEW 60

Workflows can now be created, modified, duplicated, and troubleshot from a connected MCP-enabled AI client using natural language, including inspecting run history and manually triggering workflows.

.
.
.
.

— Describes capability but no exact command or config Release Notes 2026 (2026-11-05)

WAS THIS USEFUL?  

1Password

Feedback ⓘ 1Password CLI 2.42.0 (2026-11-04) BLOG ↗

RANK

All the latest releases from the 1Password team

> 1Password's biggest change this window is on the CLI side: `op environment secret list` is now GA, 12 new Shell Plugins landed, and three existing shell plugins (Hugging Face, LocalStack, Sourcegraph) ship breaking authentication changes that require manual migration. Desktop and browser apps also picked up smaller additions — Android sharing history, Linux search and Environments UI tweaks, Mac secure-snippet formatting changes, and browser-level shortcut and API-key handling, alongside the removal of Phantom wallet saving from the browser extension.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 Breaking authentication changes in Hugging Face, LocalStack and Sourcegraph plugins

95

BREAKING

The Hugging Face shell plugin now uses the `hf` command instead of `huggingface-cli` — run `op plugin init hf` and remove the old `huggingface-cli` alias or the plugin stops injecting credentials. The LocalStack shell plugin now uses `LOCALSTACK_AUTH_TOKEN` instead of the deprecated `LOCALSTACK_API_KEY`, so setups relying on `LOCALSTACK_API_KEY` no longer receive credentials. The Sourcegraph shell plugin now supports the current `sgp_` access token format and requires the Endpoint field, since each token is tied to a specific endpoint; items missing it will stop working.

Migrate an existing Hugging Face shell plugin setup to the new `hf` command after upgrading to 2.42.0.

```
$ op plugin init hf
```

— Exact migration command and before/after given for three plugins.

1Password CLI 2.42.0 (2026-11-04)

02 `op environment secret list` command reaches GA NEW

80

The `op environment secret list` command is now generally available, letting you list all secrets in a 1Password Environment — name, value, and concealment status —

without opening the web UI.

Audit all secrets in a 1Password Environment — including which are concealed — without opening the web UI.

```
$ op environment secret list
```

— Exact command and output fields given, directly runnable.

1Password CLI 2.42.0 (2026-11-04)

---

03 **Extended credential support in Terraform, OpenTofu and Gitea plugins**

IMPROVED

70

The Terraform shell plugin now provides credentials to `terraform output` and `terraform taint` commands; the OpenTofu shell plugin now provides credentials to `tofu workspace` commands, matching Terraform; and the Gitea shell plugin now writes tea's configuration to the location tea reads when `XDG_CONFIG_HOME` is set.

— Named commands and env var but no migration steps given.

1Password CLI 2.42.0 (2026-11-04)

---

04 **12 new Shell Plugin integrations** NEW

65

1Password CLI adds Shell Plugins to authenticate the CockroachDB CLI (`cockroach sql`), Doppler CLI, DuckDB CLI with MotherDuck, LocalXpose CLI (`loclx`), Netlify CLI, npm and pnpm, Octopus Deploy CLI, OpenBao CLI (`bao`), Postman CLI, Spacelift CLI (`spacectl`), Stackhero CLI, and xAI Grok CLI via Touch ID and other unlock options.

— Names all twelve plugins but no setup command shown. 1Password CLI 2.42.0 (2026-11-04)

THINNER COVERAGE BELOW

---

05 **Toggle for keyboard shortcut hints in sign-in prompts** NEW

55

Users can turn off keyboard shortcut hints in sign-in prompts via Settings > Appearance & shortcuts in the browser extension.

Reduce visual clutter in sign-in prompts by disabling keyboard shortcut hints for all users who find them distracting.

> 📍 In the browser extension, go to Settings › Appearance & shortcuts and turn off keyboard shortcut hints.

— Exact settings path and example given. 1Password in the browser 8.12.46-26 (20…

---

06 **Removed Phantom wallet save integration** BREAKING 55

The option to save Phantom wallet items to 1Password directly from the Phantom browser extension has been removed.

— Clear before/after but no migration guidance given.

1Password in the browser 8.12.46-26 (20…

07 **Secure snippets formatting changes in Mac Labs** IMPROVED 50

In 1Password for Mac Labs, the secure snippets feature now always shows the formatting toolbar when editing a snippet, and snippets fill without formatting as plain text.

— Two concrete behavior changes but no config detail.

1Password for Mac beta 8.12.46-26 (2026…

08 **Automatic API key credential field population** NEW 50

When saving an API key from sites like GitHub or OpenAI, 1Password now automatically adds it to the API Credential item's credential field.

— Automatic behavior described, no explicit user action needed.

1Password in the browser 8.12.46-26 (20…

09 **Environments summary panel in Developer settings (Linux)** NEW 45

1Password for Linux now shows a summary explaining what 1Password Environments are under Developer > View Environments.

— Exact navigation path, but no mechanism detail.

1Password for Linux beta 8.12.46-26 (20…

10 **View Sharing History for items on Android** NEW 40

Android users can view an item's sharing history by selecting View Sharing History from the dropdown menu in item details.

— UI path given but no further mechanism detail. 1Password for Android beta 8.12.46-26 (…

11 **Recent search history in Linux search bar** NEW 40

The 1Password for Linux search bar now displays recently opened items when selected, so frequently used items are reachable without retyping.

— Clear UI behavior but no settings or limits given.

1Password for Linux beta 8.12.46-26 (20…

12 **Updated MCP Server setup instructions (Linux)** IMPROVED

40

The 1Password MCP Server setup instructions now point to Developer settings and include added troubleshooting steps.

— Docs update noted but no specifics on new steps.

1Password for Linux beta 8.12.46-26 (20…

13 **Dedicated Autofill section in Linux Settings** IMPROVED

35

Autofill-related settings have been moved into a dedicated Autofill section in 1Password for Linux's Settings.

— Simple reorg with a clear settings location. 1Password for Linux beta 8.12.46-26 (20…

14 **Watchtower share button renamed to 'Post to X'** IMPROVED

35

The Watchtower share button in 1Password for Linux was renamed from `Tweet` to `Post to X`.

— Exact rename named but a trivial cosmetic change.

1Password for Linux beta 8.12.46-26 (20…

15 **Responsive sidebar and search bar behavior (Linux)** IMPROVED

30

The 1Password for Linux sidebar now collapses automatically when the app window narrows, and the search bar centers when the window resizes.

— Describes UI behavior but no user action involved.

1Password for Linux beta 8.12.46-26 (20…

16 **Improved Autofill in Opera on Android** IMPROVED

25

1Password for Android improves Autofill support in the Opera browser.

— Thin one-line improvement with no detail on what changed.

1Password for Android beta 8.12.46-26 (…

└──▷ **BREAKING ON UPGRADE**

! The Hugging Face shell plugin now uses `hf` instead of `huggingface-cli`; existing users must run `op plugin init hf` and remove the old `huggingface-cli` alias or the plugin will not inject credentials.

! The LocalStack shell plugin now uses `LOCALSTACK_AUTH_TOKEN` instead of the deprecated `LOCALSTACK_API_KEY`; setups relying on `LOCALSTACK_API_KEY` will no longer receive credentials.

! The Sourcegraph shell plugin now requires the Endpoint field, since each `sgp_` token is tied to a specific endpoint; items missing the Endpoint field will stop working.

! The option to save Phantom wallet items to 1Password directly from the Phantom extension has been removed.

WAS THIS USEFUL? 👍 👎

Clerk Feedback ⓘ

Separate session lifetimes for native apps (2026-11-04) BLOG ↗ RANK

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

> Clerk now lets developers configure native app session lifetimes separately from web browser sessions.

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

01 **Separate session lifetimes for native apps** NEW 35

Native app sessions now support their own maximum lifetime setting, independent of web browser session lifetimes, allowing different expiration policies for native vs. web clients.

— Describes the change but no config key, UI path, or mechanism given.

Separate session lifetimes for native a…

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ Ory Network, Ory Kratos, Ory Polis v26.4.11 released

NOTES ↗ RANK

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

> Ory Polis tightens SAML assertion validation to prevent issuer spoofing.

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

01 **SAML assertion issuer validation in Ory Polis** IMPROVED

45

Ory Polis now requires the Issuer of a signed SAML assertion to match the entity ID of the IdP defined in its connection, rejecting assertions from unexpected issuers.

— Explains mechanism and scope but no config/flag to act on

Ory Network, Ory Kratos, Ory Polis v26.…

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

amneziawg-installer

Feedback ⓘ v6.2.0 NOTES ↗

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

amneziawg-installer now gives ARM64 servers the 3.1 profile via DKMS with auto-rebuild on kernel change; recovery commands corrected for prebuilt-module servers.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v6.2.0 https://github.com/bivlked/amneziawg-installer.git
# already have the repo? check out this version:
$ git checkout v6.2.0
```

NEW SUPPORTED PLATFORM ARM64 servers now get the AmneziaWG 3.1 profile by default via DKMS from the PPA, with automatic module rebuild on kernel change. Use `--protocol=2.0` on ARM64 to keep the prebuilt-package path.

The commands `manage check`, `manage repair-module`, and `install --repair-module` now all give the correct recovery instruction for ARM64 servers running a prebuilt module after a kernel update.

└──▷ BREAKING ON UPGRADE

! An ARM64 server installed with an earlier version keeps its existing generation (2.0); moving it to 3.1 requires a full reinstall with every client profile reissued.

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ Android 1.5.18 (2026-11-04) BLOG ↗ RANK

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

> Firezone's Android client gains EU control plane selection alongside parallelized WireGuard encryption for improved tunnel throughput.

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

01 **EU control plane selection in Settings** NEW 45

Android users can now select between the Firezone US or EU control plane from a redesigned Settings screen.

— Names exact UI path but no further mechanism detail Android 1.5.18 (2026-11-04)

02 **Parallelized WireGuard tunnel encryption** IMPROVED 40

Tunnel throughput now scales across worker threads via parallelized WireGuard encryption and decryption.

— Explains mechanism but gives no benchmarks or config Android 1.5.18 (2026-11-04)

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

Aembit

Feedback ⓘ

Aembit Now Secures OpenAI Dots Personal Agents (2026-11-04) 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 extended its MCP Identity Gateway to secure OpenAI Dots personal agents, adding tool-level access control, blended-identity authorization, just-in-time credential issuance, and a per-agent kill switch — all through the same connector and policy plane already used for ChatGPT, Claude, Copilot Studio, and Meta Muse.

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

01 **OpenAI Dots support via MCP Identity Gateway** NEW 70

Aembit now secures OpenAI Dots personal agents through the MCP Identity Gateway, reusing the same single connector already used for ChatGPT — no new product to deploy. The same policy framework is applied across Dots, ChatGPT, Claude, Copilot Studio, and Meta Muse from a single centralized policy control plane.

Extend existing ChatGPT gateway coverage to also control OpenAI Dots, limiting what tools the agent can call on a connected MCP server.

📍 In the Aembit console, go to the MCP Identity Gateway connector already configured for ChatGPT, extend the access policy to include the Dots agent identity, and add an MCP Tool Access Control rule specifying which tool names or patterns are permitted for `tools/call` — Dots will then share the same gateway connection with no additional connector needed.

— Names connector reuse and five covered agents; console steps given, no API detail.

Aembit Now Secures OpenAI Dots Personal…

02 **MCP Tool Access Control for agent requests** NEW 70

MCP Tool Access Control filters which tools a Dot can discover via `tools/list` and checks each `tools/call` before it reaches the MCP server, with enforcement failing closed.

Extend existing ChatGPT gateway coverage to also control OpenAI Dots, limiting what tools the agent can call on a connected MCP server.

📍 In the Aembit console, go to the MCP Identity Gateway connector already configured for ChatGPT, extend the access policy to include the Dots agent identity, and add an MCP Tool Access Control rule specifying which tool names or patterns

are permitted for `tools/call` — Dots will then share the same gateway connection with no additional connector needed.

— Names exact protocol methods and fail-closed behaviour; has console steps.

Aembit Now Secures OpenAI Dots Personal…

03 **Just-in-time scoped credentials for MCP requests** NEW

60

Aembit issues short-lived, scoped credentials at runtime for each MCP server request so agents never hold static long-lived credentials for connected services such as GitHub, Microsoft 365, and Salesforce.

— Names covered services and credential model but no setup steps.

Aembit Now Secures OpenAI Dots Personal…

THINNER COVERAGE BELOW

04 **MCP Authorization Server blended identity policy** NEW

55

The MCP Authorization Server evaluates policy on blended identity — the agent, the employee it acts for, group membership, and runtime security posture — before issuing credentials.

— Explains evaluation factors but gives no config or endpoint to act on.

Aembit Now Secures OpenAI Dots Personal…

05 **Kill switch to revoke agent access** NEW

45

Security teams can revoke a specific Dot's access through Aembit without affecting the employee's own accounts, acting as an independent kill switch for agent access.

— Describes the capability but no mechanism or revocation path detail.

Aembit Now Secures OpenAI Dots Personal…

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ November 4, 2026 NOTES ↗

RANK

DigiCert CertCentral is a web-based platform for ordering, issuing, managing, and renewing SSL/TLS certificates and other digital trust credentials.

> DigiCert CertCentral now automatically matches ICA chain key type to CSR key type for products with ICA selection enabled, and warns administrators when a manually selected chain doesn't match.

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

01 ICA chain key type matching and mismatch warnings IMPROVED

78

CertCentral now checks the CSR key type (RSA or ECC) against the selected ICA chain, reachable via `Settings > Product Settings > Allowed intermediate chains` menu. When a product's default chain is set to DigiCert default, CertCentral auto-selects a chain matching the CSR's key type, provided ICA selection is enabled and both key types are supported; this applies on new orders and renewals but not on reissues or duplicates. When a specific chain is set instead, CertCentral shows a warning on key-type mismatch (e.g. an ECC certificate with an RSA ICA chain) but still allows submission.

Verify whether ICA selection is enabled for a product and confirm its default chain before ordering an ECC certificate.

> 📍 In CertCentral, go to Settings › Product Settings. Select the product (for example, Basic OV). Look for the 'Allowed intermediate chains [Intermediate CA] › [Root CA]' menu. If present, check the Default intermediate chain menu: if set to DigiCert default (recommended), CertCentral will now auto-select a key-type-matching chain; if set to a specific chain (for example, 'DigiCert Global G3 TLS ECC SHA384 2020 CA1 (SHA384ECDSA) › DigiCert Global Root G3 (SHA384ECDSA)'), a mismatch with the CSR key type will trigger a warning.

— Names exact UI path and behavior, but single release only November 4, 2026

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Nov 3rd, 2026 NOTES ↗

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

Compyl adds an in-app feedback dialog in settings and tightens the Incidents and Contracts APIs to require the `users` field.

Adds an in-app feedback dialog in the settings menu to share feedback, report bugs, or request features with an optional impact rating and follow-up preference.

NEW INTEGRATION Vercel integration

└──▷ BREAKING ON UPGRADE

! The Incidents and Contracts APIs now require the `users` field on both create and update requests; requests that omit it will fail validation.

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

OneTrust

Feedback ⓘ

October 26, 2026 - cmpapi-202610.1.0 (2026-10-26) BLOG ↗ RANK

The OneTrust Developer Portal allows you to leverage APIs and SDKs to integrate external systems seamlessly and streamline the data flow with the OneTrust platform.

> OneTrust's Developer Portal added eight new Assessment Automation APIs with richer data on existing assessment endpoints, six Mobile App Consent APIs (four now with ETag-based conditional requests), DROP record support for California-registered data brokers, granular OAuth scopes across several API families, and bulk data-element and purpose actions in the Bulk Export API.

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

01 Mobile App Consent APIs with conditional request support NEW

95

Six new Mobile App Consent APIs retrieve banner and preference center interface attributes, active IAB and Google vendor data, and consent interactions: `Get Banner Data API`, `Get Preference Center Data API`, `Get IAB and Google Vendors API`, `Log Consent API`, `Get Universal Consent Purposes API`, and `Submit Global Opt-Out Request API`. Four of these — `Get Universal Consent Purposes API`, `Get Banner Data API`, `Get Preference Center Data API`, and `Get IAB and Google Vendors API` — now support ETag-based conditional requests via the `If-None-Match` request header and `ETag` response header, returning `304 Not Modified` when data is unchanged to reduce bandwidth consumption.

Poll the Get Banner Data API without re-downloading unchanged data — cache the ETag from the first response and send it on subsequent calls to get a 304 when nothing has changed.

```
$ # First call — capture the ETag
curl -i 'https://<your-domain>/api/v1/banner-data' \
  -H 'Authorization: Bearer <token>' \
  > response.txt

# Extract ETag value
ETAG=$(grep -i '^etag:' response.txt | awk '{print $2}' | tr -d
'\r')

# Subsequent call — skip download if data is unchanged
curl -i 'https://<your-domain>/api/v1/banner-data' \
  -H 'Authorization: Bearer <token>' \
  -H "If-None-Match: $ETAG"
```

— Named endpoints, headers, and a runnable polling example

October 26, 2026 - cmpapi-202610.1.0 (2… | October 16, 2026 - 202610.1.0 (2026-10-…

02 **Expanded data and context on existing Assessment APIs**

IMPROVED 90

The Get Assessment API now supports an `include` query parameter to retrieve additional datasets (information requests, notes, and stage history); the Get Assessment API and Get Assessment Record API now return information request metrics, flagged question counts, and the `questionResponseReviewStatus` object covering AI and human review statuses; and the Launch Assessment API now supports an `oneAgent` query parameter to launch assessments in OneAgent contexts.

Fetch assessment datasets including information requests, notes, and stage history in one call — useful when building a compliance dashboard that needs full assessment context.

```
$ curl -X GET
'https://<instance>.onetrust.com/api/assessment/v2/assessments/<a
ssessmentId>?include=informationRequests,notes,stageHistory' \
  -H 'Authorization: Bearer <token>' \
  -H 'Accept: application/json'
```

— Runnable curl example demonstrates the new parameter

October 16, 2026 - 202610.1.0 (2026-10-…

03 **Granular OAuth scopes for access management** NEW 85

Adds more granular OAuth scopes: `ORGANIZATION_READ`, `ORGANIZATION_WRITE`, and `ORGANIZATION_DELETE` on the Organizations APIs; `USER_GROUP_READ`,

`USER_GROUP_WRITE`, and `USER_GROUP_DELETE` on the User Groups and User Groups V2 APIs; `USER_READ`, `USER_WRITE`, and `USER_DELETE` on the Users V2 APIs; and `ATTACHMENT_WRITE` now accepted by the Upload File API in addition to the existing `ATTACHMENT` scope.

Request only the read permission needed to list users, reducing blast radius for a read-only integration token against the Users V2 API.

```
$ curl -X POST 'https://<instance>.onetrust.com/oauth/token' \
  -d 'grant_type=client_credentials&client_id=<id>&client_secret=
<secret>&scope=USER_READ'
```

— Scopes named with a runnable token-request example

October 16, 2026 - 202610.1.0 (2026-10-…

04 **New Assessment Automation APIs** NEW 75

Adds new assessment APIs: the Get Assessment Question Section API, Get Next Assessment Question API, and related navigation endpoints to navigate questions and retrieve unanswered items; the Perform Assessment Workflow Action API to trigger state changes on assessments programmatically; the Submit Child Question Responses API to submit child question responses; the Manage Assessment Question Attachments API to manage file uploads and downloads on assessment questions; the Get List of Assigned Assessments API and Get List of Linked Assessments by Entity ID API to retrieve assessments assigned to a user or linked to an entity; and the Search Assessments API to search across assessments.

— Endpoints named but no usage example given October 16, 2026 - 202610.1.0 (2026-10-…

05 **Bulk data-element and purpose actions in Bulk Export API** NEW 70

Expands the Create Bulk Export API with the new `DATA_ELEMENT_BULK_ACTIONS` export type to bulk-delete or rename data element values across matching data subjects, and the new `REMOVE_PURPOSE` action in the `DATA_SUBJECTS_BULK_ACTIONS` export type to remove purposes from data subjects in bulk.

— Export types named, no usage example October 16, 2026 - 202610.1.0 (2026-10-…

06 **DELETE Request and Opt-out Platform (DROP) record APIs** NEW 65

Create, update, and retrieve DELETE Request and Opt-out Platform records with the new Create Data Subject Request for DROP Record API, Update DROP Record API, and Get List of DROP Records API to support California-registered data brokers.

— Endpoints and compliance context named, no example call

October 16, 2026 - 202610.1.0 (2026-10-…

07 **Consent audit info in anonymous receipts** IMPROVED 60

Anonymous consent receipts now always include consent audit information in `dsDataElements` — returning consent model and interaction type data regardless of whether Advanced Analytics is enabled — ensuring evidence is available for auditing and compliance investigations.

— Field named but no request example or endpoint path

October 26, 2026 - cmpapi-202610.1.0 (2…

08 **Conditional requests on Get Data Subject's Preferences V2 API**

IMPROVED 60

Determine whether preference data has changed since last retrieval using the ETag and Last-Modified response headers now returned by the Get Data Subject's Preferences V2 API.

— Headers and endpoint named, no example provided

October 16, 2026 - 202610.1.0 (2026-10-…

THINNER COVERAGE BELOW

09 **Primary record selection for custom entity types** IMPROVED 55

Set primary records for `MODELS`, `AISYSTEMS`, `AIAGENTS`, `DATASETS`, and `PROJECTS` custom entity types via the updated Set Primary Record API.

— Entity types and endpoint named, no example call

October 16, 2026 - 202610.1.0 (2026-10-…

NEW COMPLIANCE FRAMEWORK GPP support for Maryland, Indiana, Kentucky, Rhode Island, and Minnesota

October 6, 2026 - cmpapi-202609.2.0 (2026-10-06) BLOG ↗ STABLE

CMP APIs add an OT-Is-Anonymous-User header to block profile creation, plus new accessibility and age-gate fields across banner, preference center, and consent log endpoints.

CMP APIs now accept the OT-Is-Anonymous-User request header; set it to `true` to mark a request as anonymous and prevent data subject profile creation even when an OT-Identifier is present.

Get Preference Center Data API and Get Banner Data API now return the `toggleDisableHints` response parameter with localized guidance for age-restricted preference toggles.

> Get Preference Center Data API now returns the `disabledToggleAriaLabel` response parameter to improve accessibility support for disabled controls via standardized ARIA labels.
>
> Get IAB and Google Vendors API, Get Banner Data API, Get Preference Center Data API, and Log Consent API now return the `closeButtonAriaLabel` response parameter inside the `appConfig` object.
>
> Get Preference Center Data API now returns `isConsentDisabled` and `isLegIntDisabled` parameters inside the `purposes` object to distinguish Age Gate-restricted purposes from other disabled consent controls.

WAS THIS USEFUL?  

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 3.0.0 (2026-11-04) 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 3.0.0 release introduces a Hint scanner action with configurable fallback handling, new MCP tools for managing agent GitHub access, Terraform resources for secrets/webhooks/notifications, and several breaking changes to run-trace and run-status APIs, alongside customer-managed web gateways, nested policy field scoping, and new Analytics views.

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

01 **Breaking changes to agent run endpoints** BREAKING 95

`get_agent_run_trace` adds an `events_after_seq` input to page through raw run events, with callers passing each page's `next_after_seq` back until `has_more` is false and the endpoint no longer returns a total event count; and while a run is still in progress, `GET /agents/{agent_id}/runs` and `GET /agents/{agent_id}/runs/timeline` now return `event_count`, token counts, and `total_cost_usd` as `null` until the run ends and usage is recorded.

— Exact param names and before/after behavior fully specified Exact endpoints and fields named, no migration guidance given

3.0.0 (2026-11-04)

02 **Hint scanner action with configurable fallback** NEW 90

A scanner can now be set to Hint to let a flagged tool call through while warning the agent the output may be untrusted; the warning is carried in `_meta["com.runlayer/hint"]` and its text is editable under Edit security messaging. Admins can also choose how clients that can't apply Mask or Hint are handled — Block (default) or Alert — via new settings under Security Scanners → Settings → Clients that can't apply Mask or Hint.

— Names meta field, UI path and settings but no code sample 3.0.0 (2026-11-04)

03 **RunlayerBlockedError on blocked Mask/Hint output** BREAKING 90

`afterTool` and `runTool` in `@runlayer/hooks-sdk` now throw `RunlayerBlockedError` when a Mask or Hint scanner would rewrite tool output and

the fallback is Block — previously the original unmasked output was passed through to the caller.

— Clear before/after with exact error and SDK functions named `3.0.0 (2026-11-04)`

04 **Connector tool exposure control in Plugin** NEW 90

Control connector tool exposure in the Runlayer Plugin by sending `X-Runlayer-Tool-Exposure: force-initial` to keep a client on the search-based flow, or toggle flat mode for self-hosted deployments with `FF_DISABLE=plugins-flat-mode` (disable) or `FF_ENABLE=plugins-flat-mode` (enable, for deployments still on WorkOS feature flags).

Keep a specific MCP client on the search-based tool flow instead of receiving connector tools directly.

```
$ X-Runlayer-Tool-Exposure: force-initial
```

— Exact header and env-var flags given with usage example `3.0.0 (2026-11-04)`

05 **Terraform resources for secrets, webhooks, notifications** NEW 80

New `runlayer_secret_store`, `runlayer_agent_webhook`, and `runlayer_agent_notification_route` Terraform resources let teams manage external secret stores, inbound webhooks, and failure notification routes as code.

— Three resource names given with what each manages `3.0.0 (2026-11-04)`

06 **Nested object field policy conditions** NEW 75

Policy conditions can read fields inside object arguments using `payload.<arg>.<field>` syntax (for example `payload.folder.id`), letting a Deny rule be scoped to a single folder.

Scope a Deny policy to writes inside a specific folder by reading a nested object field in the payload.

```yaml
payload.folder.id
```

— Named syntax with concrete example, limited to one use case `3.0.0 (2026-11-04)`

07 **GitHub access management via MCP tools** NEW 75

Agent GitHub access can now be managed from any AI client using the new MCP tools `get_agent_github_access`, `grant_agent_github_access`, and `revoke_agent_github_access`.

— All three tool names given but no usage detail beyond names 3.0.0 (2026-11-04)

THINNER COVERAGE BELOW

08 **Spend treemap in Analytics** NEW 55

The Analytics Spend tab now shows a treemap of spend broken down by department, agent, model, or user, with drill-down to entity spend details.

— Names breakdown dimensions but no exact navigation steps 3.0.0 (2026-11-04)

09 **Chrome extension auto-update from org URL** NEW 40

Organizations can turn on Chrome extension updates from their own update URL via Settings → AI Watch → Manage settings → Browser extension.

— Exact UI path given but minimal mechanism described 3.0.0 (2026-11-04)

10 **Cache read/write breakdown in Agent Monitor** IMPROVED 40

The Agent Monitor Tokens tile now shows cache reads and cache writes as separate line items.

— Minor UI change, no further mechanism or path given 3.0.0 (2026-11-04)

11 **Customer-managed web gateway for agent traffic** NEW 20

Agent internet traffic can now be routed through a customer-supplied HTTPS web gateway, applying its own filtering and logging alongside Runlayer's network policy.

— Described only in prose with no named config or endpoint 3.0.0 (2026-11-04)

└──▷ **BREAKING ON UPGRADE**

! `get_agent_run_trace` no longer returns a total count with a run's raw events; page with the new `events_after_seq` input, passing each page's `next_after_seq` back until `has_more` is false.

! While a run is still in progress, `GET /agents/{agent_id}/runs` and `GET /agents/{agent_id}/runs/timeline` return `event_count`, token counts, and `total_cost_usd` as `null` until the run ends and usage is recorded.

! `afterTool` and `runTool` in `@runlayer/hooks-sdk` now throw `RunlayerBlockedError` when a Mask or Hint scanner would rewrite tool output and the fallback is Block — previously the original unmasked output was passed through.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive

Feedback ⓘ v0.62.0 NOTES ↗ RANK

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

> ToolHive's v0.62.0 release adds Bedrock model support for Codex, a TLS listener option for the operator, and brings the virtual MCP server up to date with the final MCP 2026-07-28 specification.

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

01 **Codex for Bedrock support in llm setup** NEW 55

ToolHive can now configure Codex for Bedrock models via the `llm setup` command.

— Names exact command but no further mechanism or options. v0.62.0

02 **vMCP spec alignment and tools/call elicitation** IMPROVED 50

ToolHive's virtual MCP server now aligns with the final MCP 2026-07-28 specification, and the virtual MCP server answers `tools/call` requests with a URL elicitation.

— Names the spec version and method but no usage detail. v0.62.0

03 **TLS listener field in operator configuration** NEW 40

Adds a TLS listener field and corresponding validation to the operator configuration.

— Names the field but not its exact config key or syntax. v0.62.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Syslifters SysReptor

Feedback ⓘ 2026.102 NOTES ↗

A customizable and powerful penetration testing reporting platform for offensive security professionals. Simplify, customize, and automate your pentest reports with ease.

SysReptor adds multi-tab AI agents with background execution and history, PDF metadata via HTML meta tags in designs, and background task worker processes.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2026.102
https://github.com/Syslifters/sysreptor.git
# already have the repo? check out this version:
$ git checkout 2026.102
```

Set PDF metadata directly in report designs via HTML `<meta>` tags.

AI agent now supports multiple concurrent tabs, background execution, full session history including subagent messages, and deleting findings and notes.

AI agent sessions are automatically cancelled when a user's access is revoked.

Tasks now run in background worker processes.

NEW COMPLIANCE FRAMEWORK HTB CDEP certification report design

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

Bots - See which AI assistants act for your users in Account Abuse Protection (2026-11-0

BLOG ↗ RANK

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

> Cloudflare shipped a new Account Abuse Protection dashboard with population and per-account views plus AI assistant visibility and filtering, alongside AI Gateway model aliasing, a read-only Browser Isolation mode, session affinity for private load balancers, and zone-level bot policies for personal AI agents.

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

01 **Account Abuse Protection dashboard with population and account views**

NEW 88

Cloudflare adds a new Account Abuse Protection dashboard showing a population overview with total login and sign-up events, unique accounts, IPs, devices, top countries and ASNs, and leaked credential check results. The individual account view shows a single Hashed User ID's top networks, locations, devices, and an event log with Ray IDs, and lets you apply a Security Rule directly from that view. Two new roles are added: Account Abuse Protection grants dashboard access, and Account Abuse Protection PII additionally grants access to account-level PII (such as email addresses) and is required to create or edit Logpush jobs that include PII.

Investigate a suspicious Hashed User ID showing leaked credential events and apply a mitigation rule without leaving the dashboard.

📍 In the Cloudflare dashboard, go to Security › Account Abuse Protection. Under the Population overview, filter by leaked credential events to find high-risk accounts. Click a Hashed User ID to open its Individual Account overview, review the event log and Ray IDs, then select Mitigate to create a Security Rule scoped to that Hashed User ID.

— Detailed views, roles, and dashboard navigation all named

Bots - New Account Abuse Protection Das…

02 **AI assistant visibility and filters in Account Abuse Protection**

NEW 76

Account Abuse Protection now shows AI assistant activity in a population overview listing which AI assistants are active across your users and how many accounts use each, and surfaces per login and sign-up event whether an AI assistant made the request and which one. New AI assistant filters let you filter accounts by number of AI assistant events or number of distinct AI assistants, so you can investigate or mitigate a single account without blocking an assistant for every other user who relies on it. This visibility can also be used to tune agent settings in AI Crawler and Agent access policies.

Investigate unusual login activity on a single account that may stem from an AI assistant, without disrupting other users who rely on that assistant.

> 📍 In the Cloudflare dashboard, go to Account Abuse Protection › Account activity, open the target account, and use the AI assistant filters to filter by number of AI assistant events. Review which assistant made the login or sign-up requests, then adjust that account's policy in AI Crawler and Agent access policies without applying a global block.

— Specific filters and workflow described, lacks exact UI path

Bots - See which AI assistants act for …

---

03 **Model aliasing in AI Gateway** NEW 75

AI Gateway now accepts a single model alias to route requests across a pool of models with fallback ordering or smart routing, letting you swap backend models without redeploying your app. Configure it via Cloudflare dashboard under AI Gateway > your gateway > Model Aliases, adding models in fallback order or enabling smart routing, then calling the alias name in requests.

Configure an AI Gateway model alias to route across a pool of models with fallback ordering, so a single name in your app points to multiple backends without code changes on model swaps.

> 📍 In the Cloudflare dashboard, go to AI Gateway › your gateway › Model Aliases, create an alias name, add models in fallback order or enable smart routing, then call that alias name in your requests.

— Clear mechanism and dashboard path, no API given product docs

---

04 **Read-only keyboard mode in Browser Isolation** NEW 75

Browser Isolation adds an Allow read-only actions keyboard setting (beta) that lets users search, copy, print, and zoom on isolated websites while blocking typed input from reaching the page. Enable it in Zero Trust > Gateway > Firewall Policies > Browser Isolation by editing the relevant policy.

Let users search, copy, and zoom on an isolated website in Browser Isolation without allowing keyboard input to reach the page.

📍 In the Cloudflare dashboard, go to Zero Trust › Gateway › Firewall Policies › Browser Isolation, edit the relevant policy, and enable the Allow read-only actions keyboard setting.

— Names exact setting and dashboard path product docs

THINNER COVERAGE BELOW

05 **Session affinity and hash steering in private load balancers**

NEW

55

Private load balancers now support Hash endpoint steering and session affinity, keyed on the Cloudflare One Client device or the source IP address, to steer Cloudflare One users and private network traffic to the same endpoint.

— Mechanism described but no config key or steps given product docs

06 **Zone-level bot policies for personal AI agents** NEW

30

Bot management adds zone-level policies for top personal AI agents, allowing granular access control without applying the same action to all agent traffic.

·

— Thin one-line description with no config detail product docs

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 5, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue