THE CYBER TOOLCHAIN

NO. 142

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED NOVEMBER 17, 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 198 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 55 of 57 tools

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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

Three unrelated things stand out. Twingate's Identity Firewall brokers Kubernetes, SSH and web app sessions through an open-source

vim ? help

Gateway instead of standing credentials. Tenzir v7.0.0-rc.3 pushes queries down into ClickHouse, Splunk, Defender and SentinelOne, so detection content runs where the data already sits. Skylos turns scan findings into a planned, test-verified `cleanup` queue.

PROTECT

Grant a Kubernetes or SSH session without issuing a standing credential

Twingate's Identity Firewall puts privileged access to Kubernetes, SSH and web apps behind the open-source Twingate Gateway with Layer 7 controls, which is the difference between a user holding a kubeconfig forever and getting one brokered session. Arcade's project API keys now carry scope and can be deactivated, and OAuth connections get named profiles, so a leaked integration key no longer means every tool it touched.

Twingate · Arcade

DETECT

Run detection queries against data you never ingested

Query pushdown to ClickHouse, Splunk, Defender (KQL), SentinelOne, SQL Server, MySQL and DuckDB means a hunt can execute inside the source system rather than copying everything into a central lake first — the usual tax on cross-tool detection. The reworked OCSF functions (`ocsf_derive`, `ocsf_cast`, `ocsf_compact`, `ocsf_trim`) handle the normalization that makes results from those sources comparable.

Tenzir

PROTECT

Close a finding instead of just filing it

Skylos's `cleanup` workflow plans a fix, has an LLM review it, and verifies it against tests before it lands — the step that usually sits in a human's queue for a quarter. Sonatype adds in-UI triage for quarantined components, so a blocked build gets resolved where the block happened rather than in a side channel.

skylos · Sonatype

RESPOND

Restore a pod and have it come back configured correctly

Restore Hooks run init containers and exec commands against restored pods, and storage-class and image-repository remapping ConfigMaps let a

> restore target a different cluster shape than the one it was taken from. That is the gap between having a backup and having a recovery that actually boots in your DR environment.
>
> Velero

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

NetSPI OCInferno

Feedback ⓘ changes since 2026-11-16 CODE ↗

OCInferno adds SMTP credential enumeration, attack-path filters, and CSV/XLSX export to the network resources module

└──▷ TRY IT

Narrow an attack-path search to paths originating from users or groups and landing only at the tenancy level, cutting noise from child-compartment grants.

```
$ ocinferno process_og_attack_paths --graph-json
/path/to/oracle_cloud_hound.json --from-type user,group --to-scope
ocid1.tenancy
```

Trace all attack paths from a specific user including IDD admin compound chains, capping results at 50 paths.

```
$ ocinferno process_og_attack_paths --graph-json
/path/to/oracle_cloud_hound.json --from someuser --all-paths --max-
paths 50
```

Export network inventory results to CSV or XLSX using the new `--csv-out` and `--xlsx-out` flags on `process_network_resources`. Not checked yet

Filter attack-path starting nodes by principal type using the new `--from-type` flag (e.g. `user,group`) on the OpenGraph attack paths module. Not checked yet

Restrict attack-path targets to a specific scope using the new `--to-scope` flag (e.g. `ocid1.tenancy`) to skip child-compartment grants. Not checked yet

Enumerate IDD user SMTP credentials via the new `list_smtp_credentials` and `save_idd_smtp_credentials` methods, returning SCIM metadata including creation status and expiry. Not checked yet

Network inventory now includes load-balancer backends as a navigable endpoint kind (`KIND_BACKEND`), surfaced in the inventory report alongside IPs, FQDNs, URLs, and CIDRs. Not checked yet

+2 more

WAS THIS USEFUL? 👍 👎

adscan

Feedback ⓘ v15.0.0 NOTES ↗ RANK

> adscan v15.0.0 adds engagement scoping, a licensing system for ADscan PRO, and significantly expands what `adscan recheck` can re-verify, alongside smaller additions to reporting and scan startup.

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

01 Expanded recheck coverage for AD CS, delegation and DC settings

Not checked yet IMPROVED 90

`adscan recheck` now covers AD CS findings (ESC1–ESC9, ESC11, ESC13, ESC15), delegated rights, LAPS and gMSA password reads, DCSync rights, and DC signing/channel-binding/NTLM/LDAPS/Kerberos RC4 settings. Rechecks can run as a chosen account via `-u` / `--username` and `-p` / `--password`, target a specific attack path with `--path`, and both `adscan recheck` and `adscan engagement` can write only the JSON document to stdout with `--json`.

Re-verify AD CS misconfigurations and DC security settings for a specific attack path after remediation, using a read-only account that is not stored.

```
$ adscan recheck --path ap-42 -u auditor -p 'R3adOnly!' --json > recheck-results.json
```

— Lists every finding type and flag, with a runnable command example v15.0.0

02 Engagement scoping by host and account Not checked yet NEW

88

Adds `--scope-allow` and `--scope-deny` on `adscan ci` and `adscan start`, accepting IPs, CIDRs, ranges, hostnames, SIDs, DNs, `DOMAIN\name`, and UPNs, plus `--scope-allow-file` and `--scope-deny-file` for file-based lists and `--scope-dc-access full` to permit full DC contact beyond directory reads. An interactive `scope` shell command lets operators inspect, change, and explain the current engagement scope, and prompts before any widening of scope.

— Names every flag and the interactive command, no example attached v15.0.0

03 License management for ADscan PRO Not checked yet NEW

84

Adds the `adscan license` command with `activate`, `status`, `check`, `deactivate`, and `relay import` subcommands. PRO can be installed with `adscan install --license-file`, and offline or air-gapped licence refresh is supported via

a file exported with `--relay-out` and imported with `adscan license relay import`.

— Full subcommand and flag set named, no runnable example given v15.0.0

---

04 **PRO licensing enforcement and distribution changes**

Not checked yet BREAKING 65

From v15.0.0, `adscan start`, `adscan ci`, `adscan execute`, and `adscan recheck` stop with an error on unlicensed machines, with `adscan ci` exiting 78. PRO images are no longer published to Docker Hub and GHCR, with new versions published only to ADscan's own registry after a transition period, and `--offline` no longer suppresses domain registration telemetry for PRO licences.

— Names exact commands, exit code and registries, no migration steps v15.0.0

---

05 **Compliance framework opt-out and policy agreement tracking**

Not checked yet NEW 64

`--frameworks none` (or selecting None) omits compliance framework mapping from audit reports while recording the assessor's decision. A per-workspace `agreed_policy` key in scan config makes both reports name the policies in effect and their digests when a run deviates from the agreed policy.

— Names flag and config key but no worked example v15.0.0

THINNER COVERAGE BELOW

---

06 **Skip minimum memory check at startup** Not checked yet NEW 47

Passing `--allow-low-memory` starts the scan without enforcing the minimum memory check.

— Single flag, thin description of the check itself v15.0.0

└──▷ **BREAKING ON UPGRADE**

! ADscan PRO now requires an activated licence; from v15.0.0 `adscan start`, `adscan ci`, `adscan execute`, and `adscan recheck` stop with an error on unlicensed machines, and `adscan ci` exits 78.

! PRO images are no longer published to Docker Hub and GHCR; new versions are published only to ADscan's own registry after a transition period.

! `--offline` no longer suppresses domain registration telemetry for PRO licences.

WAS THIS USEFUL? 👍 👎

Rapid7 Metasploit Pro

Feedback ⓘ

November 16, 2026: 5.1.0-2026111301 NOTES ↗ RANK

Metasploit Pro is a penetration testing platform that automates vulnerability exploitation, network scanning, and security assessments to help professionals identify and remediate weaknesses.

> Metasploit Pro 5.1.0 adds 16 new exploit and auxiliary modules spanning Citrix NetScaler, GitLab CI/CD, Portainer, OpenBullet2, Flowise, SPIP and HashiCorp Consul, alongside a new Linux ELF injector persistence module, Windows support for Docker Image Persistence, expanded payload architecture coverage, and shared Zimbra SOAP infrastructure.

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

01 **New exploit modules for Citrix, GitLab, Portainer, OpenBullet2, Flowise, SPIP, Consul**

Not checked yet NEW 85

Adds an unauthenticated RCE via command injection for `Citrix NetScaler` (CVE-2026-88771), reported exploited in the wild prior to disclosure; an integer overflow in `GitLab` CE/EE Linux amd64's CI/CD regular expression compiler enabling RCE as the git user (CVE-2026-97229); a `Portainer` exploit that uses valid credentials to create a privileged container, bind-mount the host filesystem, and execute a payload as root on the Docker host; two `OpenBullet2` modules — an RCE via plain C# execution mode (CVE-2026-25856) and a credential disclosure scanner that forces NTLMv2 hash disclosure via a malicious UNC path (CVE-2026-39908); an account takeover auxiliary module for `Flowise` versions 3.0.5 and earlier (CVE-2025-58434) that retrieves a valid password reset token directly from the API response; an unauthorized password change auxiliary module for `SPIP` before 4.4.23 via a validly-signed POST without `id_auteur` authorization checks; and a `HashiCorp Consul` version scanner.

— every CVE and mechanism named, but no exact module path given

November 16, 2026: 5.1.0-2026111301

02 **Linux ELF injector persistence module** Not checked yet NEW

64

New persistence module injects into Linux ELF binaries using code cave and segment injection techniques, supporting `x86`, `x64`, and `AArch64` architectures with architecture-specific trampolines that preserve the original execution flow.

— mechanism and supported architectures named, no exact command shown

November 16, 2026: 5.1.0-2026111301

03 **Windows support in Docker Image Persistence module**

Not checked yet IMPROVED 64

The Docker Image Persistence module now supports Windows, enabling Meterpreter sessions on Windows hosts running Docker Desktop to create a persistent Alpine container with access to the host's C: drive.

— mechanism and host requirement named, no command shown

November 16, 2026: 5.1.0-2026111301

THINNER COVERAGE BELOW

04 **Payload support for loongarch64 and AArch64 Windows**

Not checked yet NEW 48

Adds `loongarch64` shell payloads and the required configs to work with fetch payloads, and adds an AArch64 Windows DLL template, hooking up payload generation to use it.

— names new payload targets but no usage walkthrough November 16, 2026: 5.1.0-2026111301

05 **Shared Zimbra SOAP support across modules** Not checked yet NEW

48

Adds shared Zimbra SOAP support for account and administrator authentication, account discovery, and version retrieval, letting Zimbra modules reuse consistent token routing, response validation, and error handling.

— describes shared infrastructure but no specific module or endpoint named

November 16, 2026: 5.1.0-2026111301

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

SpecterOps BloodHound

Feedback ⓘ v9.10.0 NOTES ↗

RANK

BloodHound is an Active Directory reconnaissance tool that maps attack paths and privilege escalation routes within Windows domain environments.

> BloodHound's v9.10.0 release expands its API with endpoints for relationship and node kind lookups, management operation control, and client artifact storage, alongside new S3-backed artifact storage and admin UI actions for support bundles.

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

01 **Relationship kind and node kind lookup endpoints** Not checked yet

NEW 70

Adds `api/v2/relationship-kinds/{relationship_kind_id}` to retrieve details about a specific relationship kind, and a new GET node kind endpoint to look up node kind details by ID.

Retrieve details about a specific relationship kind by its ID to understand edge types in your graph.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://<bloodhound-host>/api/v2/relationship-kinds/42'
```

— Exact endpoint path and curl example given; node kind path not specified. v9.10.0

02 **Client artifact download, delete and S3 storage** Not checked yet

NEW 65

Adds a download client artifact endpoint and a delete client artifact endpoint to retrieve and remove client artifacts via the API. Adds `part_number` support to the artifact upload flow for multipart uploads, and S3 configuration support to store and retrieve artifacts from S3-compatible object storage.

— Names config/flow surfaces but no exact endpoint paths or config keys. v9.10.0

THINNER COVERAGE BELOW

03 **Management operations control and inspection endpoints**

Not checked yet NEW 55

Adds queue management operation endpoints to control and inspect running management operations via the API, plus a get-available management operation endpoint to query which management operations are available. Artifact metadata is now exposed on management operations responses.

— Describes operation types but omits exact endpoint paths. v9.10.0

---

04 **Download and delete client support bundle in admin UI**

Not checked yet NEW 45

Adds a download client support bundle UI action and a delete client support bundle UI action to the administration interface.

— UI action named but no exact navigation path given. v9.10.0

WAS THIS USEFUL? 👍 👎

◆ Password & Credential Cracking

brutespray

Feedback ⓘ v2.7.6 NOTES ↗ RANK

BruteSpray is a Python tool that automatically performs credential brute-forcing against services discovered by Nmap scans using Medusa.

> BruteSpray v2.7.6 adds telnet decoy-shell detection to catch fake logins and fixes circuit-breaker logic so `-max-conn-fails` works without requiring spray mode.

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

01 **Circuit breaker fix for `-max-conn-fails`** Not checked yet

IMPROVED 84

Passing `-max-conn-fails` alone now enables the circuit breaker under `-skip-policy auto` without also requiring `-spray` or an explicit policy flag, allowing host skipping outside spray mode. Hosts mixing connection errors and auth rejections now stop after the configured `-max-conn-fails` threshold instead of spraying to the end of the credential list, a behaviour change versus v2.7.5.

Skip a host that keeps rejecting connections during a non-spray scan without needing to set an explicit spray policy.

```
$ brutespray -f nmap.gnmap -u admin -p passlist.txt -max-conn-fails 3
```

▶ View command output v2.7.6

— Named flags plus a runnable example command v2.7.6

02 **Telnet decoy-shell detection** Not checked yet IMPROVED 69

After an apparent successful telnet login, BruteSpray now sends a confirmation probe; if a second authentication screen appears, the result is downgraded from a success to a failed credential and a `WARN` `telnet-second-screen` finding is emitted. This prevents fake-shell logins on DVRs, IP cameras, and managed switches from polluting success files. Telnet targets previously reported as successful under this condition now report the new finding instead, so old success files from telnet runs should be re-checked.

— Explains mechanism and before/after but no CLI flag to toggle it v2.7.6

└──▷ BREAKING ON UPGRADE

! Hosts that mix connection errors and auth rejections now stop after the configured `-max-conn-fails` threshold instead of spraying to the end of the credential list — expected behaviour change versus v2.7.5.

! Telnet targets previously reported as successful when a decoy shell was present now report `telnet-second-screen` findings instead; re-check old success files from telnet runs against this pattern.

WAS THIS USEFUL?  

APPLICATION SECURITY

◆ Web Application Security

Invicti

Feedback ⓘ changelog · Nov 17, 2026 CHANGELOG ↗

Official documentation for Invicti AppSec Platform, Invicti Platform, and Invicti Enterprise — web application security testing, DAST, API security, and vulnerability management.

Invicti adds target-type filtering on the Scans page, configurable default Accept headers, GraphQL schema visibility in the API catalog, removed-operation tracking, API gateway connection testing, NTA renaming without token regeneration, and REST API support for linking APIs from a URL.

Filter scans by target type (Web or API) and see the target type in a new column on the Scans page. Not checked yet

Set the default `Accept` header the scanner sends on web and API requests using the `Default Web Accept Header` and `Default API Accept Header` settings in a target's advanced settings. Not checked yet

Link an API from a URL through the Invicti Platform REST API, matching what was previously only possible in the UI. Not checked yet

UPDATED DETECTION The API catalog now shows GraphQL schemas with the same detail it shows for REST APIs. Not checked yet

The API catalog now marks operations that no longer appear in a re-imported API definition as `Removed`, with a status filter to show or hide them. Not checked yet

+7 more

UPDATED DETECTION Stateful API scanning now covers DELETE endpoints

UPDATED INTEGRATION SCIM-provisioned users assignable to vulnerabilities and secrets

└─▷ BREAKING ON UPGRADE

! Email notifications for new and updated APIs are now off by default — teams that relied on receiving these automatically must re-enable them in profile notification settings.

WAS THIS USEFUL? 👍 👎

◆ Code Security

GitHub Feedback ⓘ

Dependabot alert dismissal requests now last one year (2026-11-17) BLOG ↗

RANK

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

> GitHub's biggest change this window is making the Microsoft Copilot harness the default (and eventually only) VS Code agent runtime, replacing the Local harness ahead of a February 2027 removal; alongside that it shipped new SKU-level and all-AI-credits budget controls, a `strictPluginOnlyCustomization` managed setting for locking down Copilot customization sources, and a broad set of weekly Copilot updates covering GHES mobile support, VS Code Agents window UI, new models, dictation languages, and accessibility settings.

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

01 Copilot harness becomes default agent runtime in VS Code

Not checked yet 95

BREAKING

The Microsoft Copilot harness is now the default agent session runtime in VS Code (selectable via Copilot under Session Target), built on the Copilot SDK and shared with Copilot CLI and the GitHub Copilot app, replacing the Local harness which will be removed in February 2027 along with the `vscode.temporaryDefaultToLocalHarness` enterprise-managed setting — existing customizations are not automatically converted and must be reviewed for compatibility. The harness adds Copilot Fusion, which automatically chooses models and, on supported tiers, coordinates drafting and review for a task; assisted permissions that reduce approval interruptions while still surfacing calls that fail risk assessment for manual confirmation; Copilot connectors PREVIEW to bring in context from team services without configuring MCP servers manually; canvases PREVIEW for planning work and exploring interactive diagrams on task boards; continuity of the same agent session across the editor and the Agents window while the session's Agent Host process stays running; and administrator-enforced managed runtime permissions and sandbox restrictions that the Local harness could not enforce.

Switch an existing VS Code Copilot agent session from the deprecated Local harness to the new Copilot harness before the February 2027 removal deadline.

📍 In VS Code, start a new agent session and under Session Target select Copilot.

— Exact migration steps, named setting, multiple concrete capabilities

GitHub Copilot weekly releases — Novemb… | VS Code retires the Local harness in Fe…

02 strictPluginOnlyCustomization managed setting for Copilot

Not checked yet | NEW | 95

Copilot customizations can be restricted to plugins and managed settings — blocking standalone user and project configs — with the `strictPluginOnlyCustomization` managed setting, applicable to any combination of skills, custom agents, hooks, MCP servers, and executable Copilot extensions. Pairing it with `strictKnownMarketplaces` also restricts which marketplaces supply eligible plugins, letting enterprise admins distribute approved MCP tools and canvases through a shared plugin library.

Lock MCP servers and executable extensions to plugin-only sources while leaving skills, agents, and hooks open to per-user configuration.

json

```
{
  "strictPluginOnlyCustomization": {
    "mcpServers": true,
    "extensions": true,
    "skills": false,
    "customAgents": false,
    "hooks": false
  }
}
```

Combine plugin-only enforcement with marketplace restrictions so only approved marketplace sources can supply eligible plugins.

json

```
{
  "strictPluginOnlyCustomization": {
    "mcpServers": true,
    "extensions": true
  },
  "strictKnownMarketplaces": true
}
```

— Exact config keys and JSON examples provided | Control where GitHub Copilot customizat…

03 SKU-level and all-AI-credits budget controls

Not checked yet | NEW | 78

GitHub billing adds SKU-level budgets to cap overage spending on a specific product SKU, such as Copilot code review, Secret Protection AI credits, or GitHub Advanced

Security — including selecting Advanced Security and the SKU for AI Scan and agentic autofix — and an `All AI credits` budget that caps total overage spending across every SKU that consumes AI credits (Copilot, Copilot cloud agent, Spark, and Advanced Security AI features), drawing on the shared AI credits allowance before counting overage.

Cap AI credit overage across all GitHub AI features org-wide to prevent unexpected spend from Copilot, Spark, and Advanced Security AI.

> 📍 In GitHub billing settings, go to Budgets and create a new budget, select `All AI credits` as the budget scope, and set a spending limit. Available where overage spend is enabled.

— Clear UI steps and named SKUs, no API endpoint given `product docs`

---

04 **Agents window enhancements in VS Code** `Not checked yet` **NEW**

75

The VS Code Agents window adds Inbox mode to the sessions list to track agent sessions across all workspaces from one view, surfacing sessions needing attention first and letting you dismiss, hold, or snooze the rest; lets you move a chat out of the Agents window into its own window **PREVIEW** to arrange the chat, its editors, and its terminals independently; and lets you find, install, and uninstall extensions without leaving the Agents window, showing a warning for extensions with unsupported features.

Use Inbox mode in VS Code to triage all agent sessions across workspaces and act on the ones waiting for input.

> 📍 In VS Code, open the Agents window, switch to the sessions list, and enable Inbox mode — sessions needing attention appear first; use dismiss, hold, or snooze on the rest.

— Detailed behaviour with usage steps for Inbox mode

`GitHub Copilot weekly releases — Novemb…`

05 **Compare Models agent tool in VS Code** Not checked yet NEW

70

Agent-driven model comparisons can be enabled in VS Code by turning on the Compare Models agent tools in settings, letting agents find available models, launch comparisons, and summarize findings.

Let an agent benchmark available models against each other and surface a summary, without running comparisons manually.

> 📍 In VS Code settings, enable Compare Models under agent tools, then open the Agents window and ask the agent to compare available models for your task.

— Named setting with usage steps GitHub Copilot weekly releases — Novemb…

06 **New Copilot models: Claude Fable 5.2 and Gemini 4 Carbon**

Not checked yet NEW

60

Adds `Claude Fable 5.2` and Gemini 4 Carbon models, now available to Copilot Pro+, Max, Business, and Enterprise users; Copilot Business and Enterprise admins must enable `Claude Fable 5.2` via a policy that is off by default.

— Names models and policy but no exact setting path

GitHub Copilot weekly releases — Novemb…

THINNER COVERAGE BELOW

07 **GitHub Mobile enhancements: cross-device sessions and iPhone Duo support**

Not checked yet NEW

55

GitHub Mobile now lets you continue cloud and Codespace sessions on a phone via a new Open in menu with a GitHub Mobile option showing a QR code, and adds support for iPhone Duo with a layout that adapts when the device is unfolded and vertical tabs for navigating between Copilot, Pull Requests, Issues, and Notifications.

— Names menu but no deeper mechanism Describes behaviour but no setup steps

GitHub Copilot weekly releases — Novemb… Take GitHub Mobile with you on iPhone D…

08 **High contrast accessibility settings replace Increase contrast**

Not checked yet

55

BREAKING

High contrast can now be combined with red, green, or blue-yellow color filters in `Settings > Accessibility`, with a preview of combinations before applying; High contrast replaces the separate Increase contrast setting.

— Named setting path, notes removal of old setting

GitHub Copilot weekly releases — Novemb…

09 **Zero data retention exemption for Claude models** Not checked yet

NEW

50

Claude Fable 5, 5.1, and 5.2 offer zero data retention (ZDR) through 2026 via a time-bound exemption; enterprises must contact GitHub to request access, and approved enterprises have requests routed automatically to the ZDR endpoint once an administrator enables those models.

— Names models and policy but process is vague product docs

10 **Custom field filtering for issues and pull requests**

Not checked yet NEW

50

Issues and pull requests can be filtered by custom fields, such as a GitHub Project's Size, from the Add filter menu.

— Names exact menu but little elaboration GitHub Copilot weekly releases — Novemb…

11 **Budget limit raised to 40,000 per account** Not checked yet

NEW SUPPORTED PLATFORM IMPROVED

45

GitHub Enterprise Cloud accounts can now create up to 40,000 budgets per account, up from 10,000, applying automatically on github.com and GHE.com with nothing to turn on.

— Simple numeric limit change, no further mechanism

Budget limit increased to 40,000 budget…

12 **MCP server reauthentication without leaving conversation**

Not checked yet IMPROVED

45

Users can sign back in to MCP servers such as Slack without leaving the conversation — when a server needs re-authentication, the conversation offers the matching connector and resumes the task.

— Explains behaviour, no exact command GitHub Copilot weekly releases — Novemb…

13 **Multilingual dictation in Copilot** Not checked yet NEW

45

Dictate in French, Spanish, and other languages by downloading the multilingual voice model in settings and choosing Auto or a specific language.

— Clear steps but limited scope detail GitHub Copilot weekly releases — Novemb…

14 **GHES support in Copilot mobile app, GA** Not checked yet NEW

40

The Copilot app now supports working with repositories, issues, and pull requests from GitHub Enterprise Server, generally available.

— States GA status but no setup detail GitHub Copilot weekly releases — Novemb…

15 **Scheduled Project automations down to the minute** Not checked yet

IMPROVED 40

Project automations can now be scheduled down to the minute, or using a custom cron schedule such as weekdays at 9:30.

— Example given but no exact UI path GitHub Copilot weekly releases — Novemb…

16 **Dependabot alert dismissal requests last one year** Not checked yet

IMPROVED 35

Dependabot delegated alert dismissal requests now stay open for one year before expiring, up from one week, giving reviewers more time to respond without requiring resubmission.

— Simple timing change, no new surface to act on Dependabot alert dismissal requests now…

└──▷ **BREAKING ON UPGRADE**

! The Local harness will be removed in February 2027; workflows must be migrated to the Copilot harness before then.

! High contrast replaces the separate Increase contrast setting.

! The `vscode.temporaryDefaultToLocalHarness` enterprise-managed setting will be removed in February 2027 along with the Local harness; keeping the exception configured will not keep the Local harness available.

! The Local harness is removed in February 2027; existing customizations are not automatically converted when switching to the Copilot harness and must be reviewed for compatibility.

> Bulk start AI Scan for repositories (2026-11-16) BLOG ↗ STABLE
>
> GitHub now allows bulk-starting AI Scans for up to 20 repositories at once from the code scanning organization alerts page, with no Copilot license required.
>
> Start AI Scans for up to 20 repositories at once from the code scanning organization alerts page or the organization-level setting, without requiring a GitHub Copilot license. Not checked yet
>
> Start an AI Scan for a single repository directly from its alerts page, with findings appearing in the code scanning backlog alongside CodeQL alerts. Not checked yet

WAS THIS USEFUL? 👍 👎

syzkaller Feedback ⓘ changes since 2026-11-13 CODE ↗

Syzkaller is an unsupervised coverage-guided kernel fuzzer developed by Google that automatically finds bugs in OS kernels by generating and executing random syscall sequences.

syzkaller adds automatic AI cause bisection scheduling, KCSAN flow improvements, and hardware breakpoint delay injection for race analysis

Adds a hardware breakpoint delay helper to `race_toolkit.h` that opens a `PERF_TYPE_BREAKPOINT` on a pid/address and attaches a `BPF_PROG_TYPE_PERF_EVENT` BPF program to stall inside the `#DB` handler for a configurable number of microseconds, aiding precise data-race reproduction. Not checked yet

Adds `FailureDetectableBy` output field to the KCSAN assessment flow to surface which layer can observe a detected data race. Not checked yet

Adds `open_tree_attr` syscall descriptions to `sys/linux`. Not checked yet

Adds x86 SYZOS `memwrite` instruction support and expands SMRAM coverage in the executor. Not checked yet

Adds automatic AI cause bisection scheduling to the dashboard, including rate limiting for older bugs via `bisectCauseNewBugAge`, `maxBacklogBisectCauseJobs`, and `backlogBisectCauseWindow` constants to prevent old bugs from saturating agents, and a `maxAutoBisectCauseAttempts` cap to bound automatic retries of failed or aborted cause-bisection workflows. Not checked yet

+5 more

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ changelog · Nov 16, 2026 BLOG ↗ RANK

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

> Snyk added browser-based login support for authenticated API scanning of Postman and Bruno targets, removing the need for manual token management on OAuth 2.0, SAML, and SSO-protected APIs.

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

01 **Browser login for Postman and Bruno API targets** Not checked yet

NEW 81

A new Browser login option under Settings > Authentication lets Snyk replay a browser-based sign-in at scan time and automatically extract the session credential, enabling authenticated scans against API targets that don't issue static API keys. It supports OAuth 2.0 Authorization Code, SAML, and SSO flows, including two-factor authentication (2FA), and works with scheduled scans. The authentication settings page now groups options into Collection-defined login, Browser login, and Logout detection sections for Postman and Bruno targets.

Schedule an authenticated DAST scan against a Postman or Bruno collection that uses OAuth 2.0, SAML, or SSO — without pre-generating or rotating tokens manually.

> 📍 In the Snyk Web UI, open your Postman or Bruno API target, go to Settings › Authentication, select Browser login, configure the sign-in flow, and save. Snyk will replay the browser-based login automatically at each scheduled scan.

— Names exact UI path, flows, and 2FA support; no limits given 2026-11-16-8b960aa6

WAS THIS USEFUL? 👍 👎

Mend.io Mend SAST Feedback ⓘ

Version 26.10.2 (16-November-2026) NOTES ↗

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

Mend SAST adds a guided onboarding wizard with repository selection, pilot scans, and team email delegation, plus engine-generation language filtering in Global Scan Configuration.

Start SAST scanning for repositories from the Mend Platform using a new guided onboarding wizard that handles repository selection and pilot scan execution without manual configuration. Not checked yet

Delegate SAST onboarding to colleagues by sending email invitations from within the wizard to drive team-wide adoption. Not checked yet

Filter Global Scan Configuration language options by engine generation so only supported languages appear during scan setup. Not checked yet

UPDATED DETECTION JavaScript Gen 2 engine — DOM XSS detection via jQuery AJAX callbacks

UPDATED DETECTION Swift Gen 2 engine — false-positive reduction for CWE-117 and CWE-295

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ DAST updates (2026-11-16) BLOG ↗ RANK

Veracode is a cloud-based application security platform that scans code for vulnerabilities using static, dynamic, and software composition analysis testing methods.

> Veracode's latest DAST updates focus on finer-grained target configuration, expanded Enterprise-mode authentication and reporting, and new ways to plug DAST into the broader Veracode Platform and automation via REST APIs.

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

01 Expanded DAST target configuration options

Not checked yet

IMPROVED

70

New DAST target configuration options include Directory Restrictions to control which URL paths the scanner may follow, a business unit selection when creating a web application or API specification target (or on the DAST configuration page), and a custom host-to-IP address mapping for when DNS lookup is unavailable or undesired. Configuration also gains an optional Domain field for HTTP basic authentication, the ability to assign multiple teams to a single target, and Full scan Enterprise mode as the new default, with the legacy Full scan option deprecated and no longer available for new or Quick scan targets.

Authenticate to a Windows domain-integrated web application in a DAST scan by supplying domain credentials alongside the username and password.

> 📍 In the DAST target configuration, go to Authentication › HTTP Basic Authentication, fill in Username and Password, then set the optional Domain field to your Windows domain name (e.g. CORP).

Restrict a DAST scan to a specific subdirectory of your application so the crawler does not stray into unrelated paths.

> 📍 In the DAST target configuration, go to Allowed URLs, add the URL, then enable Directory Restrictions and enter the permitted path prefix (e.g. `/app/v2/`).

— Names several config fields with steps; no deeper mechanism given

DAST updates (2026-11-16)

02 New authentication methods in DAST Enterprise mode

Not checked yet NEW

63

DAST Enterprise mode adds AI-assisted login for complex login flows, login scripts, scanner variables, client certificates for web applications, Scriptable Request Modifications (SRM) for APIs, time-based one-time password (TOTP) seeds (including

via the REST API), and custom cookies. API Scanning also adds the ability to use JavaScript to modify an HTTP request at runtime when authenticating with a remote host (Scriptable Request Modification).

— Names all auth surfaces but no configuration steps shown DAST updates (2026-11-16)

THINNER COVERAGE BELOW

03 **ISM proxy exclusions and EU region support** Not checked yet

IMPROVED 55

ISM endpoint `24.10.0` adds a proxy exclusion list so specified hosts bypass the configured proxy settings, and ISM is now available for Veracode Dynamic Analysis of web applications and API specifications in the European region.

— Names versioned endpoint and region but no config steps DAST updates (2026-11-16)

04 **Public REST API for DAST target and analysis management**

Not checked yet NEW 55

A Public REST API for DAST target configuration now enables automation of DAST target management and analysis triggering, and the Dynamic Analysis REST API can be used to set URL scan settings for all analyses and scans at the organization level.

— Names the API surfaces but no endpoint paths given DAST updates (2026-11-16)

05 **Open port scanning for web applications** Not checked yet NEW

45

Dynamic Analysis and DAST Enterprise mode now scan web applications for open ports, adding new detections for HTTP headers, TLS, and fingerprinting.

— Explains what it detects but no usage steps given DAST updates (2026-11-16)

06 **DAST application linking to Veracode Platform** Not checked yet

NEW 45

DAST adds application linking, letting you link analysis run results to a Veracode Platform application to evaluate against policy and combine results from multiple scan types in a single report.

— Describes the capability but no steps or endpoint named DAST updates (2026-11-16)

07 **EASM-to-DAST automatic target creation** Not checked yet NEW

40

EASM integration now automatically creates discovered targets in DAST, which can be mapped or converted to regular DAST targets for scanning.

— Explains workflow but no UI path or config given DAST updates (2026-11-16)

08 **Postman Collection scanning in API Scanning** Not checked yet NEW

40

API scanning now supports uploading and scanning Postman Collections in the Veracode Platform.

– Names the format and platform but no steps shown DAST updates (2026-11-16)

09 **DAST REST API scheduling endpoints changed** Not checked yet

40

BREAKING

The DAST REST API scheduling endpoints have changed; integrations that use the REST API to manage a schedule must be updated to use the new endpoints, per the DAST API specification.

– Flags a breaking change but doesn't name the new endpoints DAST updates (2026-11-16)

10 **Downloadable Enterprise scan reports** Not checked yet NEW

30

DAST Enterprise mode adds downloadable scan reports that provide visibility into scan coverage and scanner behavior.

– Brief description with no detail on report format or access DAST updates (2026-11-16)

NEW DETECTION Malicious package detection

NEW SUPPORTED PLATFORM Python 3.12 support in SCA agent-based scans

NEW SUPPORTED PLATFORM Poetry package manager support in SCA agent-based scans

NEW SUPPORTED PLATFORM Poetry package manager support in SCA Upload scans

└──▷ BREAKING ON UPGRADE

! The DAST REST API scheduling endpoints have changed; integrations that use the REST API to manage a schedule must be updated to use the new endpoints (see the DAST API specification for details).

SCA updates (2026-11-16) BLOG ↗ LATEST STABLE

Veracode SCA adds `--appname` and `--show-cli` scan flags, Unified Policy assignment for EU and US Federal workspaces, SBOM generation, upgrade-path-aware build-break rules, and updated backend API URLs.

Link SCA agent scan results to an application profile by passing `--appname` to the scan command (requires OAuth or HMAC authentication, Commercial region only). Not checked yet

Save JSON scan output to a file while displaying a human-readable summary in the console by passing `--show-cli` together with `--json` to the SCA agent scan command; using `--show-cli` without `--json` returns an error. Not checked yet

> Assign policies to SCA agent-based scan workspaces in the US Federal region using the Unified Policy feature to create issues and break builds based on defined criteria. Not checked yet
>
> Assign policies to SCA agent-based scan workspaces in the European region using the Unified Policy feature to create issues and break builds based on defined criteria. Not checked yet
>
> Control when builds are broken based on upgrade-path availability by configuring the Vulnerability Severity rule in a policy (Commercial region). Not checked yet
>
> +5 more

WAS THIS USEFUL?  

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.167.11 NOTES ↗ RANK

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

> Infisical v0.167.11 adds a new dynamic secret provider for OpenRouter, introduces output masking in PAM, extends the database schema with TOTP config support, and fixes the ServiceNow integration to read all workflow approvers.

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

01 **TOTP config table in database schema** Not checked yet NEW

45

Adds a `totp_configs` table to the sanitized database schema to support TOTP configuration.

— Names exact schema table but no behavior explained. v0.167.11

02 **OpenRouter dynamic secret provider** Not checked yet NEW

40

Adds a new dynamic secret provider for OpenRouter API keys, allowing Infisical to generate and manage dynamic credentials for OpenRouter.

— Names the provider but no setup or config detail given. v0.167.11

03 **ServiceNow integration reads all approvers** Not checked yet

IMPROVED

40

The ServiceNow integration now reads every approver from the approvals column, rather than only the first one, fixing incomplete approval processing.

— Explains before/after behavior but no config surface named. v0.167.11

04 **Output masking for PAM** Not checked yet NEW

10

Adds output masking for PAM (Privileged Access Management), though the release notes do not detail the mechanism or scope.

— Mentioned only in summary with no further detail. v0.167.11

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

LavaMoat

Feedback ⓘ harden-v0.9.0 NOTES ↗ RANK

LavaMoat is a security tool that uses JavaScript sandboxing to protect Node.js and browser apps from supply chain attacks by isolating package permissions.

> LavaMoat's harden-v0.9.0 release refines its runner to better handle bin script conflicts and package manager detection, aimed at blocking more classes of bin script confusion.

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

01 **Bin script and package manager handling in harden runner**

Not checked yet IMPROVED 35

LavaMoat harden's runner now accepts new bins that are not collisions with existing entries, reducing false positives in bin script confusion detection. It also falls back on the detected package manager more eagerly when no explicit manager is configured.

— Describes behavior change but no exact flags, configs, or commands given. harden-v0.9.0

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ

Manage organizations with the new Organizations APIs (2026-11-03) BLOG ↗

RANK

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

> Fastly introduced a new Organizations API that lets administrators programmatically manage organization-level accounts, member invitations, memberships, and security requirements like SSO and two-factor authentication.

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

01 **Organizations API for accounts, members, and security**

Not checked yet NEW 58

The new Organizations API adds endpoints to get and update an organization, including requiring SSO or two-factor authentication for every member. It also covers managing accounts (list existing accounts, create new ones under the organization), member invitations (invite people, list invitations, resend an invitation, cancel an invitation), and memberships (list members filtered by name, email address, or role, change a member's roles, and remove members).

– Describes endpoint capabilities in prose but no exact paths or methods given

Manage organizations with the new Organ…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.16.9 NOTES ↗ CODE ↗ RANK

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

> Netcap shipped exotic link-layer decoding with explicit device profiling, hard memory bounds across TCP, reassembly, and aggregate capture state, and an expanded MCP server with paginated, count-only inventory queries.

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

01 **Exotic link-layer decoding and device profiles** Not checked yet

NEW 95

Adds explicit device profile mapping for exotic link types via the new `-link-profiles` flag on `net capture` (accepted in PCAP, live, and service modes); reconstructs link-layer payloads for ALP, MS/TP, 6LoBAC, and ROHC frames with source ranges and transformations preserved in audit records; decodes Darwin USB capture headers and adds bounded `LinkFrame` paths for AX.25, MS/TP, Bluetooth H4, I2C, MPEG-TS, Upper PDU, ALP, and MDB link types; and extracts BACnet NPDU/APDU metadata, MPEG-TS program/PES and H.264/H.265/AAC headers, Bluetooth ACL/L2CAP/ATT/SMP metadata, and USB descriptor correlations in upper-layer profiles.

Capture from a PCAP containing exotic link types (e.g. MPEG-TS or Bluetooth H4) and supply an explicit device-profile mapping so Netcap knows how to frame each link layer.

```
$ ./net capture -read exotic_traffic.pcap -out nc-exotic -link-profiles ./my-device-profiles.json
```

— Names every protocol and flag with a runnable example. v0.16.9

02 **Bounded memory limits for TCP, reassembly, and aggregate state**

Not checked yet NEW 95

Sets default bounds of 64 MiB for `--max-retained-tcp-bytes` and 32 MiB for `--max-retained-reassembly-bytes`, aborting capture rather than growing without limit when exceeded, to back TCP control/reassembly payloads, normalized Software observations, vulnerability/exploit deduplication, and optional network-layer payloads with private bounded storage. Separately, `--max-aggregate-state-bytes` (default 32 MiB, also readable from env var `NC_MAX_AGGREGATE_STATE_BYTES`) caps resident Connection, Host, DeviceProfile, Service, and UDP cache state and spills excess to private scratch storage, with `0` keeping everything in memory.

Live capture with a hard ceiling on TCP reassembly memory — useful on memory-constrained sensors where runaway reassembly must abort rather than exhaust RAM.

```
$ sudo ./net capture -iface eth0 --max-retained-reassembly-bytes 16777216 --max-retained-tcp-bytes 33554432 -out nc-live
```

Limit resident aggregate cache to 64 MiB when analysing a large PCAP to keep peak RSS manageable; excess Connection/Host/Service state spills to scratch storage automatically.

```
$ ./net capture -read traffic.pcap -out nc --max-aggregate-state-bytes 67108864
```

Use the environment variable form to set the aggregate state cap in a containerised or scripted deployment without modifying the command line.

```
$ NC_MAX_AGGREGATE_STATE_BYTES=67108864 ./net capture -read traffic.pcap -out nc
```

— Flags, env var and defaults all named with runnable commands. v0.16.9 v0.16.8

03 **MCP server inventory querying and paging** Not checked yet NEW

75

Netcap's `mcp` command exposes current analysis and operational workflows through its stdio Model Context Protocol server, described as a 119-tool MCP surface; it supports count-only inventory queries that scan all matching audit records without transferring record bodies, and paginates inventories with byte-bounded pages, stable generations, search, matched totals, generation checks, structured field/JSON Pointer projection, and complete oversized JSON artifacts.

— Lists concrete capabilities but no example command given. v0.16.8 v0.16.9

THINNER COVERAGE BELOW

04 **Automatic tshark follower namespace detection** Not checked yet

IMPROVED

50

Discover the installed `tshark` follower namespace automatically — using modern `tls.stream` or falling back to legacy `tcp.stream` after an unsupported-field diagnostic.

— Names the fields involved but no usage example. v0.16.9

05 **PPI/DLT 192 wireless capture decoding** Not checked yet NEW

35

Netcap now decodes PPI/DLT 192 capture headers and preserves wireless FCS provenance.

— Names the format but gives no further mechanism or example. v0.16.8

06 **Signed Windows CLI binaries** Not checked yet IMPROVED

35

Sign production Windows CLI executables before archive and checksum generation; snapshot builds remain unsigned.

— Clear but thin operational note, no reader action. v0.16.9

07 **Stricter automatic Protobuf identification** Not checked yet

IMPROVED

30

Require multi-field protocol evidence for automatic Protobuf identification while preserving configured single-field decoding.

— States the behavior change but no configuration detail. v0.16.9

WAS THIS USEFUL? 👍 👎

sniffnet

Feedback ⓘ v1.6.0 NOTES ↗

Sniffnet is an open-source network monitoring tool that lets users easily track and analyse their internet traffic in real time.

Sniffnet adds live data rate display, a 20-category service taxonomy, an expanded view mode, and PCAP export path visibility.

└──▷ GET THIS VERSION ▾

Homebrew

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

Docker

```
$ docker pull ghcr.io/gyulyvgc/sniffnet:v1.6.0
```

git

```
$ git clone --branch v1.6.0 https://github.com/GyulyVGC/sniffnet.git
# already have the repo? check out this version:
$ git checkout v1.6.0
```

Homebrew lists this version · published under the project's GitHub account · checked 17 Nov

Wrong command?

Sniffnet now shows current data rates in the UI so you can see live throughput at a glance. Not checked yet

Sniffnet now classifies upper-layer services into 20 categories via a service taxonomy view. Not checked yet

Sniffnet adds an expanded view setting that slims the header and hides the footer to maximize traffic data on screen. Not checked yet

Show the output file path in the Overview page when exporting a PCAP file. Not checked yet

Show more details about dropped packets in the UI. Not checked yet

NEW SUPPORTED PLATFORM IPFIX collector support

NEW SUPPORTED PLATFORM IGMP protocol support

NEW SUPPORTED PLATFORM VLAN-tagged connection support

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Lacework FortiCNAPP

Feedback ⓘ changelog · Nov 17, 2026

CHANGELOG ↗ RANK

Lacework FortiCNAPP is a cloud-native application protection platform that provides security monitoring, threat detection, and compliance across cloud infrastructure and workloads.

> Lacework FortiCNAPP added a new exclusion filter condition for accounts in Explorer.

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

01 'Is none of' account filter in Explorer

Not checked yet NEW

55

Explorer queries now support an `is none of` condition for accounts, letting users exclude specific accounts from query results to narrow scope by exclusion rather than inclusion.

Exclude specific AWS accounts from an Explorer query so results show only accounts you care about.

> 📍 In Explorer, open the account filter, select `is none of`, and enter the account IDs you want to exclude from the query results.

— Names exact filter condition and UI path, lacks broader mechanism detail.

changelog-20261117-7cf2d4a1

WAS THIS USEFUL? 👍 👎

Prowler

Feedback ⓘ 5.50.0 NOTES ↗ CODE ↗ RANK

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

> Prowler 5.50.0 adds a confidence-tuning key to cut secret-scanning false positives, three new AWS checks covering Bedrock guardrails, invocation job encryption, and CloudWatch data protection, seven new AI Security Framework requirements, and tightens session handling so SAML SSO and API key tokens stay scoped to their issuing organization.

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

01 New AWS checks for Bedrock and CloudWatch data protection

Not checked yet NEW 78

Three new AWS checks ship in this release: `bedrock_guardrail_iam_condition_enforcement_enabled` verifies that IAM policies allowing `bedrock:InvokeModel` or `bedrock:InvokeModelWithResponseStream` enforce a guardrail via the `bedrock:GuardrailIdentifier` condition key; `bedrock_model_invocation_job_output_encrypted_with_cmk` verifies that S3 output from Bedrock model invocation jobs is encrypted with a customer-managed KMS key (CMK); and `cloudwatch_log_group_data_protection_policy_configured` verifies that CloudWatch log groups have a data protection policy that audits and masks sensitive data such as personal information and credentials.

Run only the new Bedrock guardrail IAM enforcement check against your AWS account to verify guardrails are required on model invocation policies.

```
$ prowler aws --check
bedrock_guardrail_iam_condition_enforcement_enabled
```

— Three named checks with mechanism, one with a runnable command 5.50.0

02 secrets_confidence tuning key for secret scanning

Not checked yet NEW 75

Tune secret-scanning noise across AWS, Cloudflare, and OpenStack providers by setting `secrets_confidence` to `low`, `medium`, or `high` in the provider config. Setting it to `medium` or higher suppresses the generic password, secret, and API key rules to reduce false positives.

Reduce false positives in secret scanning by raising the minimum confidence threshold to `medium`, which silences generic password, secret, and API key rules.

```yaml
secrets_confidence: medium
```

— Named config key with exact values and example YAML 5.50.0

03 **Org-scoped sessions for SAML SSO and API keys** Not checked yet

70

BREAKING

SAML SSO and API key sessions are now scoped to their issuing organization — `POST /api/v1/tokens/switch` rejects API keys and will not move such sessions to another organization.

— Names exact endpoint and behavior change, no example call shown 5.50.0

THINNER COVERAGE BELOW

04 **AI Security Framework gains manual requirements** Not checked yet

NEW

40

The AWS AI Security Framework gains seven Manual requirements covering Bedrock data retention, retrieval-augmented generation (RAG) authorization, and agent governance.

— Describes scope but no commands or exact requirement names 5.50.0

WAS THIS USEFUL?  

◆ CNAPP

Plerion

Feedback ⓘ Load balancers behind a Global Accelerator:

NOTES ↗ RANK

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

> Plerion's latest release adds visibility into continuous scanning status via its integrations API and CLI, ships new checks that detect load balancers and AppSync GraphQL APIs exposed through AWS Global Accelerator and live accounts, and surfaces the tools used by Google ADK and Dialogflow CX agents.

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

01 Continuous scanning visibility in integrations API and CLI

Not checked yet NEW 86

The list integrations API now returns `continuousScanningEnabled` on each integration and can be filtered with the `integrationIds` parameter. The `plerion integrations list` CLI command gains a matching `--integration-ids` flag to filter to specific accounts, and its output now includes a `CONTINUOUS SCANNING` column.

Confirm continuous scanning is active for a specific AWS account integration without fetching all integrations.

```
$ plerion integrations list --integration-ids <integration-id>
```

— Names exact API field, CLI flag and output column.

Load balancers behind a Global Accelera…

02 Global Accelerator exposure checks for load balancers

Not checked yet 83

BREAKING

New check `PLERION-AWS-571` detects Application Load Balancers that are publicly accessible via an AWS Global Accelerator, tracing the attack path through the accelerator, its listener, and its endpoint group. New check `PLERION-AWS-862` detects Network Load Balancers publicly accessible via an AWS Global Accelerator the same way. Load balancers that previously passed these checks can now fail, producing new findings on the next scan.

— Named checks with attack-path mechanism and stated breaking impact.

Load balancers behind a Global Accelera…

03 AppSync GraphQL API scanning expanded

Not checked yet · NEW

77

New check `PLERION-AWS-1253` detects AppSync GraphQL APIs that accept API key authorization on live AWS accounts. New check `PLERION-AWS-1254` detects AppSync GraphQL APIs with introspection enabled on live AWS accounts. Existing checks `PLERION-AWS-525` (public AppSync GraphQL API with no WAF web ACL) and `PLERION-AWS-521` (AppSync GraphQL API with field-level logging turned off) now also run against live AWS accounts, not only against templates.

— All four check IDs and their scope change named.

Load balancers behind a Global Accelera…

04 Tool visibility for ADK and Dialogflow CX agents

Not checked yet

NEW SUPPORTED PLATFORM · NEW VULNERABILITY CHECK · NEW

70

A new Tools tab on each Google ADK agent (on Vertex AI Agent Engine and Amazon Bedrock AgentCore) lists the agent's functions, sub-agents, and built-in tools such as Google Search. A new Tools tab on each Dialogflow CX agent lists OpenAPI operations, functions, data stores, webhooks, and playbook code actions, alongside the host and sign-in method for service-calling tools.

— UI location named but no exact navigation path given.

Load balancers behind a Global Accelera…

└──▷ **BREAKING ON UPGRADE**

! Load balancers behind an AWS Global Accelerator that previously passed checks `PLERION-AWS-571` and `PLERION-AWS-862` can now fail — expect new findings on the next scan.

> Agent security now covers your GCP agents: · NOTES ↗ · STABLE
>
> Plerion adds GCP agent security, including a new workload type for Vertex AI Agent Engine with staging-bucket access via a `PlerionAccess: Granted` tag.
>
> Tag a staging bucket with `PlerionAccess: Granted` to enable vulnerability, sensitive data, and SBOM scanning for Vertex AI Agent Engine agents deployed from a staging bucket. Not checked yet
>
> Plerion now scans Vertex AI Agent Engine agents as a new GCP workload type, covering vulnerabilities, sensitive data, and SBOM — container-image-deployed agents require no setup. Not checked yet

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore Grype

Feedback ⓘ v0.122.0 NOTES ↗ CODE ↗ RANK

Grype is an open-source vulnerability scanner that detects security vulnerabilities in container images and filesystems by matching packages against known CVE databases.

> Grype adds Docker Hardened Images OSV advisory feed as a new vulnerability data source.

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

01 **Docker Hardened Images OSV advisory feed** Not checked yet NEW

35

Grype now ingests the Docker Hardened Images (DHI) OSV advisory feed as a vulnerability data source.

— Names the feed but no mechanism or usage detail v0.122.0

WAS THIS USEFUL? 👍 👎

adm-controller

Feedback ⓘ v1.39.0 NOTES ↗ RANK

> adm-controller's v1.39.0 release adds CEL-to-Wasm policy compilation and configmap access for policy servers, extends the `kwctl annotate` CLI, and raises the minimum supported Kubernetes version to 1.31.

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

01 Force and inline flags for kwctl annotate

Not checked yet

IMPROVED

80

The `kwctl annotate` command gains `--force` and `-a` flags, letting users force-overwrite existing annotations on a policy artifact or add an annotation inline without editing a separate metadata file, e.g. `kwctl annotate my-policy.wasm --force -a 'io.kubewarden.policy.author=alice@example.com'`.

Force-overwrite existing annotations on a policy artifact, or add an annotation inline, without editing a separate metadata file.

```
$ kwctl annotate my-policy.wasm --force -a
'io.kubewarden.policy.author=alice@example.com'
```

— Exact flags and a runnable example given. v1.39.0

THINNER COVERAGE BELOW

02 Minimum Kubernetes version raised to 1.31

Not checked yet

50

BREAKING

The minimum supported Kubernetes version was raised from 1.27 to 1.31.

— Before/after versions named but no migration steps. v1.39.0

03 CEL-to-Wasm policy compilation

Not checked yet NEW

35

Adds support for compiling and running CEL (Common Expression Language) policies to Wasm.

— States capability but no usage detail given. v1.39.0

04 Configmap access for policy servers

Not checked yet NEW

25

Policy servers can now access configmaps.

— Bare statement, no mechanism or scope given. v1.39.0

└──▷ **BREAKING ON UPGRADE**

! Minimum supported Kubernetes version raised from 1.27 to 1.31.

WAS THIS USEFUL? 👍 👎

◆ **SaaS Security Posture**

Nudge Security

Feedback ⓘ

Discover AI agents built on the OpenAI Agents API NOTES ↗

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

NEW INTEGRATION OpenAI Agents API integration

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

Genie Code CLI is now generally available (2026-11-18) BLOG ↗ RANK

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

> Databricks shipped a breaking change to streaming watermark behavior, GA'd the Genie Code CLI, and added governed tag propagation for classic compute, native IP address/CIDR functions, Unity Catalog agent skill sharing via OpenSharing and Marketplace, an in-workspace support form, and daily budget thresholds for Unity Gateway.

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

01 Real-time watermark propagation for streaming queries

Not checked yet BREAKING 93

New streaming queries in real-time mode now use real-time watermark propagation by default, advancing the event-time watermark within a batch so late events are filtered against the live watermark and state is evicted sooner than before. Queries restarted from an existing checkpoint keep their original behavior, but new queries behave differently; disable the change by setting `spark.databricks.streaming.realTimeMode.realTimeWatermark.enabled` to `false` before starting a new query.

Preserve the previous watermark behavior for a new real-time streaming query that must not evict state sooner than before.

env

```
spark.databricks.streaming.realTimeMode.realTimeWatermark.enabled
= false
```

— Exact config flag, behavior change and migration note all given

Databricks Runtime 19: November 16, 202…

02 Get support form in workspace

Not checked yet NEW 75

Open a Databricks support case directly from your workspace via Help > Get support form, which pre-fills issue details, auto-attaches diagnostic files, and suggests an AI-generated solution for review before you submit.

When you hit an error in your workspace and want to open a support case without leaving Databricks or manually gathering diagnostics

> 📍 In the workspace, navigate to Help › Get support, fill in the issue details (the form auto-attaches diagnostic files), review the AI-generated solution suggestion, then click Submit.

— Exact UI navigation path provided as a usage example

Submit support cases from your workspac…

03 **Daily budget thresholds for Unity Gateway** Not checked yet NEW

70

Unity Gateway budgets (including Genie) now support daily per-user thresholds in addition to or instead of monthly thresholds, with separate alert, webhook, and block-usage actions configurable per period; daily thresholds reset at 00:00 UTC.

— Names mechanism and reset time but no exact config path

Budgets now support daily thresholds in…

04 **Unity Catalog agent skills sharing via OpenSharing and Marketplace**

Not checked yet NEW INTEGRATION NEW

60

Share Unity Catalog agent skills through OpenSharing shares and Databricks Marketplace listings BETA, giving Databricks-to-Databricks recipients read-only access to each shared skill in their catalog with asynchronous provider updates.

— Describes mechanism and access model but no setup steps

Share Unity Catalog agent skills with O…

THINNER COVERAGE BELOW

05 **Governed tag propagation for classic compute** Not checked yet

NEW

55

Governed tag propagation BETA now supports classic compute running Databricks Runtime 19 or above, so tags on a source propagate to derived tables and views on both classic and serverless compute.

— Names runtime requirement but no config steps shown

Governed tag propagation supports class…

06 **IP functions for IP address and CIDR processing** Not checked yet

NEW

45

New IP functions let workspaces query, join, and process IP addresses and CIDR blocks natively; they will soon be available by default for workspaces with the compliance security profile enabled.

— No function names or query examples given IP functions will soon be available by …

07 Genie Code CLI general availability Not checked yet NEW

32

Genie Code CLI is now generally available, sharing a single Genie Code budget and free monthly usage allowance with the workspace.

— Only states GA status and shared budget, no commands

Genie Code CLI is now generally availab…

NEW INTEGRATION PostgreSQL integrated CDC pipeline (Beta)

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! New streaming queries in real-time mode now use real-time watermark propagation by default; state is evicted sooner than before. Queries restarted from an existing checkpoint keep their original behavior, but new queries will behave differently. Set `spark.databricks.streaming.realTimeMode.realTimeWatermark.enabled` to `false` before starting a new query to restore the previous behavior.

> PostgreSQL integrated CDC pipeline is now in Beta (2026-11-15) BLOG ↗
>
> STABLE
>
> Databricks adds a single integrated CDC pipeline for PostgreSQL, requiring no separate ingestion gateway, now in Beta.

> External access to views through the Iceberg REST catalog is in Beta (2026-11-13)
>
> BLOG ↗ STABLE
>
> Databricks adds Beta support for querying Unity Catalog views from external Iceberg engines via the Iceberg REST catalog.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Flashpoint

Feedback ⓘ announced · Nov 17, 2026 PRESS ↗ RANK

Flashpoint empowers the world's leading organizations with the best data for threat intelligence to protect people, places, and assets.

> Flashpoint is introducing Agentic AI to Flashpoint Ignite, anchored by a new Agentic Hub environment and five purpose-built agents that automate investigation, alerting, prioritization, and reporting workflows, slated for availability in early 2027.

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

01 **Alert Setup Agent** Not checked yet NEW 65

Converts plain-language monitoring requirements into configured alert rules, letting teams preview results before saving. Usable in Flashpoint Ignite by opening Agentic Hub, selecting the Alert Setup Agent, entering a request like 'Notify me when threat actors discuss exploits targeting financial sector VPNs', reviewing the previewed rules, then saving.

Use the Alert Setup Agent to quickly turn a plain-language brief into live alert rules without manually building query logic.

📍 In Flashpoint Ignite, open Agentic Hub, select the Alert Setup Agent, type your monitoring requirement in plain language (e.g. 'Notify me when threat actors discuss exploits targeting financial sector VPNs'), review the previewed alert rules, then save.

— Exact navigation and workflow steps provided launch-20261117-8eb6334c

02 **PIR Agent** Not checked yet NEW 65

Transforms plain-language intelligence requirements into structured Priority Intelligence Requirements (PIRs) for team review and approval, while surfacing emerging threats not covered by existing requirements. Usable via Agentic Hub by selecting the PIR Agent, entering requirements in plain language, reviewing/approving the drafted PIRs, and checking surfaced coverage gaps.

Use the PIR Agent to convert a team's intelligence priorities into structured PIRs and discover gaps in existing coverage before the next planning cycle.

📍 In Flashpoint Ignite, open Agentic Hub, select the PIR Agent, enter your intelligence requirements in plain language, review the structured PIRs the agent drafts, approve or edit them, and check the emerging-threat gaps it surfaces that your current PIRs do not address.

— Exact navigation and workflow steps provided launch-20261117-8eb6334c

THINNER COVERAGE BELOW

03 **Chat agent for multi-step investigations** Not checked yet NEW

55

A Chat agent inside Agentic Hub turns a single natural-language question into a multi-step investigation, following leads across vulnerability, credential, fraud, and threat actor intelligence and returning source-linked findings.

— Mechanism described but no concrete command or path launch-20261117-8eb6334c

04 **Intelligence Briefing Agent** Not checked yet NEW 55

Generates scheduled intelligence briefings tailored to executives, analysts, and technical teams, drawing on each team's selected alert rules, PIRs, investigations, assets, and credential exposure data.

— Clear data sources named but no usage steps given launch-20261117-8eb6334c

05 **Echosec Reporting Agent** Not checked yet NEW 55

Transforms saved Echosec searches into scheduled intelligence reports containing assessments, cited findings, and recommended actions, preserving reports over time to track how threats evolve.

— Mechanism and output described but no usage steps given launch-20261117-8eb6334c

06 **Alert Triage Agent** Not checked yet NEW 50

Pre-filters low-context alert results, flags alerts that do not match monitoring intent, and labels and archives remaining alerts, keeping a log of the actions taken.

— Behavior described but no runnable steps given launch-20261117-8eb6334c

07 **Agentic Hub environment in Flashpoint Ignite** Not checked yet

NEW 45

A new environment inside Flashpoint Ignite that brings together a chat interface and purpose-built agents for multi-step threat investigations spanning cyber, vulnerability, fraud, and identity intelligence. Announced as part of Flashpoint's Agentic AI launch, available early 2027.

— Describes scope but no concrete usage yet; not yet available launch-20261117-8eb6334c

WAS THIS USEFUL?  

isMalicious

Feedback ⓘ updated · Nov 17, 2026 NOTES ↗ RANK

IsMalicious is a threat intelligence platform that checks IPs, domains, URLs, emails, phones and crypto wallets against indexed threat indicators.

> isMalicious added access to third-party YARA and Sigma rule collections via new API routes, with quota-exempt browsing and downloads.

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

01 **YARA and Sigma rule collections via `/rules` API** Not checked yet

NEW 83

New `/rules` API routes expose third-party YARA and Sigma rules, each annotated with its source, author, and declared licence. The `yara-rules` and `sigma-rules` collections can be browsed or polled without consuming request quota, as downloads and polling are quota-exempt; access is available on Pro, Enterprise, and feed trial plans.

— Names exact routes, collections and plan access; no usage example given

snapshot-20261117

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.261116.0 NOTES ↗ RANK

OpenCTI is an open-source threat intelligence platform that helps organizations manage, analyze, and share cyber threat intelligence data.

> OpenCTI 7.261116.0 adds asset classification fields, STIX 2.0 export via the Python client API for all SDOs and SCOs, and several knowledge-management and UI refinements spanning investigation graphs, playbooks, custom fields, and config presets.

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

01 **Asset classification fields on system objects** Not checked yet

NEW 65

Added `category`, `subcategory`, `exposure`, `criticality`, `environment`, `owner`, and `platform` fields to system objects, letting teams classify and prioritize assets in their inventory.

— Names all seven new fields but not where they're set 7.261116.0

02 **STIX 2.0 export via Python client API** Not checked yet NEW

60

Export STIX 2.0 objects via the Python client API for all SDOs and SCOs, including Malware.

— Names the API surface and object scope clearly 7.261116.0

03 **Config-preset export/import for exclusion lists and taxonomies**

Not checked yet IMPROVED 60

Export and import exclusion lists via the config-preset system alongside kill chain phases, vocabularies, status templates, and case templates.

— Names all five config-preset object types covered 7.261116.0

THINNER COVERAGE BELOW

04 **Investigation graph enrichment and entity actions** Not checked yet

NEW 50

Trigger enrichment connectors directly from an investigation graph without leaving the investigation context, and use a new Add entities action on empty investigation graphs to start building from a blank canvas.

— Describes two UI actions but not full mechanism 7.261116.0

05 **Report type manipulation in knowledge playbooks** Not checked yet

IMPROVED

50

Manipulate knowledge playbooks now support adding, replacing, and removing report types on entities.

— Clear capability, no exact playbook step named 7.261116.0

06 **Decay rules and decimal type for custom fields** Not checked yet

NEW

45

Select indicator decay rules on custom fields and add a decimal custom field type.

— Names the field capability but lacks setup detail 7.261116.0

07 **Time series breakdown enabled by default on dashboards**

Not checked yet IMPROVED

40

The time series breakdown feature on dashboards is now enabled by default, with its feature flag removed.

— States before/after but no usage detail 7.261116.0

08 **Platform theme import via partial import** Not checked yet NEW

40

Import platform themes via the partial import from section menus.

— Names the UI path but not mechanism 7.261116.0

09 **Metrics XTM Hub button in UI** Not checked yet NEW

30

Added a metrics XTM Hub button in the OpenCTI UI.

— Bare mention with no location or function detail 7.261116.0

WAS THIS USEFUL?  

◆ SOAR

Torq

Feedback ⓘ changelog · Nov 17, 2026 CHANGELOG ↗ RANK

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

> Torq introduced a SOC value dashboard quantifying automation ROI and reorganized its org and workspace settings layout.

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

01 **SOC value dashboard** Not checked yet NEW 65

A new dashboard shows alerts handled, cases resolved, analyst time saved, cost savings, and capacity gained. It is found under Org Administration > Usage, letting teams review fully automated vs. analyst-handled cases and benchmark time and cost savings.

Review SOC automation ROI — see fully automated vs. analyst-handled cases and benchmark time and cost savings.

> 📍 In Torq, go to Org Administration › Usage to view the SOC value dashboard.

— Names exact metrics and navigation path changelog-20261117-42c5e4c2

THINNER COVERAGE BELOW

02 **Reorganized org and workspace settings layout** Not checked yet

IMPROVED 20

Org and workspace settings now follow a reorganized layout intended to make each setting easier to find.

— No specifics on what changed in layout changelog-20261117-42c5e4c2

WAS THIS USEFUL? 👍 👎

Tines

Feedback ⓘ Better control over Workbench limits (2026-11-16)

BLOG ↗ RANK

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

> Tines shipped a single update giving admins control over AI credit allocations and limits for Tines-provisioned AI providers.

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

01 **AI credit allocations and limits in AI Settings** Not checked yet

NEW 60

Admins can now set AI credit allocations and limits for Tines-provisioned AI providers via the AI credit allocations & limits section of the tenant's AI Settings page, allowing per-provider or per-Story caps on AI spend to prevent runaway credit consumption.

Cap AI spend per Story or provider to prevent runaway credit consumption across your tenant.

> 📍 In the Tines console, go to AI Settings › AI credit allocations & limits and set per-provider or per-Story credit limits.

— Names exact UI path but lacks mechanism or scope detail

Better control over Workbench limits (2…

WAS THIS USEFUL? 👍 👎

◆ AI SOC

Exaforce

Feedback ⓘ

Meet ExaU: Your AI Security Team, Always On The Clock | NOTES ↗ | RANK

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

> Exaforce launched ExaU, a customizable AI agent workforce that lets practitioners build, schedule, and share their own Exabots with persistent memory, voice interaction, in-console action-taking, and autonomous email/Slack outreach, all gated by approval prompts and audit logging and reasoning over the same real-time security knowledge graph as existing Exabots.

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

01 **Persistent agent memory with inspectable Memory tab**

Not checked yet | NEW | 80

Each agent exposes a Memory tab split into three sections — Lessons, Team notes, and Only you — letting practitioners inspect what the agent has retained across sessions; agents persistently retain decisions, exceptions, and past instructions so context carries over without re-explanation.

Inspect what a deployed Exabot has learned — your preferences, shared team notes, and on-the-job lessons — before trusting it with a privileged access review.

📍 In the Exaforce console, open the agent, then select the Memory tab. Review the three sections: Lessons (things the agent learned on the job), Team notes (shared across all users of this agent), and Only you (your personal preferences, visible only to you).

— Names exact tab sections with a navigation example launch-20261116-7603e723

Meet ExaU: Your AI Security Team, Alway…

02 **Custom Exabot builder with scheduled roles** Not checked yet | NEW

72

Practitioners can build custom Exabots from a template or a plain-language description of the job, each with a named role such as threat hunter, access reviewer, policy officer, briefing agent, or security chief of staff, then deploy them to run on schedules inside the Exaforce platform with no new credentials or logins to provision.

Delegate a recurring CISO briefing task to an agent so it drafts the Monday status report automatically without manual prompting each week.

> 📍 In the Exaforce console, create a new Exabot, describe its role as 'Draft the Monday morning status report for the CISO, pulling open risks and owner status from the team', set a weekly schedule, then save and deploy the agent.

— Names roles and build methods with a worked deployment example launch-20261116-7603e723

Meet ExaU: Your AI Security Team, Alway…

03 **In-console agent actions with approval and audit logging**

Not checked yet NEW 70

During a 1:1 chat session, agents can open Exaforce console pages, read them, and interact — click, type, scroll — using their own cursor; any data-changing action waits for an explicit Allow or Deny from the user. Every agent action is recorded in the Exaforce audit log, and agents operate within permission boundaries set by the deploying user.

— Describes mechanism and guardrails but no exact command

Meet ExaU: Your AI Security Team, Alway…

04 **Autonomous email and Slack actions with approval gating**

Not checked yet NEW 60

ExaU agents act autonomously over email and Slack for investigation and outreach, but any action affecting employees, systems, or communications requires approval first, and every step is recorded in the Exaforce audit log.

— Names channels and approval mechanism, no usage example launch-20261116-7603e723

THINNER COVERAGE BELOW

05 **Agent sharing and cross-agent collaboration** Not checked yet NEW

58

Practitioners can share a read-only copy of any agent conversation with specific people or the whole organization while keeping the live chat private. Agents can also pass findings and tasks to other Exabots for multi-agent handoffs without manual coordination, and agents can be shared across the security organization so, for example, an analyst's threat-hunting agent and a CISO's briefing agent work in parallel with findings rolled into a single assessment.

— Covers three sharing mechanisms but no exact UI steps launch-20261116-7603e723

Meet ExaU: Your AI Security Team, Alway…

06 **Real-time knowledge graph reasoning for ExaU agents**

Not checked yet NEW 50

ExaU agents reason over the same real-time security knowledge graph already used by Exabots, covering identities, cloud activity, devices, SaaS applications, and security telemetry, with no additional provisioning required.

— Names data domains covered but no usage path launch-20261116-7603e723

07 **Self-scheduled follow-ups and routine limits** Not checked yet

NEW

45

Agents can set their own follow-ups and run routines that end on a specific date or after a set number of runs.

— States behaviour but no concrete limit or example

Meet ExaU: Your AI Security Team, Alway…

08 **Voice and dictation mode for agents** Not checked yet NEW

30

Agents support voice mode and dictation mode so practitioners can interact hands-free while working.

— Bare feature name, no mechanism given Meet ExaU: Your AI Security Team, Alway…

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

Repeat AI Scanner attacks to measure response consistency (2026-11-16) BLOG ↗

RANK

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

> Trend Vision One's AI Scanner can now repeat each attack multiple times within a single scan to measure how consistently a target responds, configurable via config file, CLI flag, or the scan wizard.

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

01 Repeat attack passes in AI Scanner

Not checked yet NEW

95

Set `settings.repeat_count` (0–50, requires config version 2.17.0 or later) to run every attack N+1 times in a single scan, with results grouped by attack configuration across passes. The repeat count can be overridden for a single scan without editing the config file using the `--repeat` flag on `./tmas aiscan llm` (`--repeat 0` forces a single pass even if `settings.repeat_count` is higher), or set via the interactive scan wizard under Performance → Repeat count. The global `settings.max_scan_duration_minutes` limit applies across all passes, so it should be scaled up with the repeat count.

Run each attack three times in a single scan to spot inconsistent LLM responses without touching the config file.

```
$ ./tmas aiscan llm -c config.yaml --repeat 2
```

Force a single-pass scan on the fly to override a non-zero `settings.repeat_count` already set in the config file.

```
$ ./tmas aiscan llm -c config.yaml --repeat 0
```

— Names exact config keys, flag, limits, and wizard path with runnable commands.

Repeat AI Scanner attacks to measure re…

WAS THIS USEFUL?  

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability & Exposure Management

Future Architect Vuls

Feedback ⓘ v0.43.0 NOTES ↗

CODE ↗ RANK

Vuls is an open-source, agentless vulnerability scanner for Linux and FreeBSD servers that detects security vulnerabilities in installed software and packages.

> Vuls v0.43.0 adds a flag to skip redundant Trivy database refreshes and improves vendor attribution for Windows package scans.

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

01 **Skip Trivy DB update flag** Not checked yet NEW 83

Pass `-trivy-skip-db-update` to `vuls report` or `vuls tui` to skip the Trivy vulnerability database update check and download, using the DB already present in the `-trivy-cachedb-dir` path. Useful when managing the DB separately or scanning offline.

Run a report without re-downloading the Trivy DB when you manage the DB yourself or are working offline.

```
$ vuls report -trivy-skip-db-update -trivy-cachedb-dir
/path/to/trivy-cache
```

— Names exact flags and a runnable command v0.43.0

THINNER COVERAGE BELOW

02 **Windows package vendor attribution** Not checked yet IMPROVED 50

Vuls now fills Windows package vendor fields from registry output passed to `ParseInstalledPkgs`, improving vendor attribution on Windows scans.

— Names the function changed but no usage example v0.43.0

WAS THIS USEFUL? 👍 👎

defectdojo

Feedback ⓘ 3.5.200 NOTES ↗

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

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.5.200 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.5.200
```

UPDATED INTEGRATION Trivy, Grype, and Syft: package URL and CPE recording on dependency locations

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Fleet

Feedback ⓘ fleet-v4.95.0 NOTES ↗

Fleet is an open-source platform for managing and querying device fleets, enabling security teams to monitor endpoints using osquery at scale.

Fleet adds an MCP server at `/api/v1/fleet/mcp`, Google conditional access device-ID endpoints, AI-powered config profile generation, macOS patch deadlines, multi-VPP-token support, and AI tool software inventory.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch fleet-v4.95.0
https://github.com/fleetdm/fleet.git
# already have the repo? check out this version:
$ git checkout fleet-v4.95.0
```

Access a built-in MCP server (Fleet Premium) at `/api/v1/fleet/mcp`, on by default and disabled with `server_settings.mcp_disabled`, so AI tools can connect using an API-only user token. Not checked yet

NEW INTEGRATION Bind or remove an iOS, iPadOS, or Android host's Google device ID with the new `PUT /api/v1/fleet/hosts/:id/google_device_id` and `DELETE /api/v1/fleet/hosts/:id/google_device_id` endpoints (Fleet Premium) for Google conditional access. Not checked yet

Use new API endpoints that generate a configuration profile, or its name and description, from plain-language instructions using AI. Not checked yet

Add a Deadline option for macOS Fleet-maintained app patch policies so Fleet Desktop notifies end users 1 hour and 5 minutes before the deadline, then closes and updates the app (requires Fleet Desktop 1.5.0). Not checked yet

Configure patch deadline days and time per fleet and for Unassigned in Software > Library > Settings, the API, and GitOps; deadlines use each host's local time and default to Tuesday at 17:00. Not checked yet

+6 more

NEW SUPPORTED PLATFORM Parrot OS Linux platform support

NEW SUPPORTED PLATFORM Omarchy Linux LUKS disk encryption key escrow

UPDATED VULNERABILITY CHECK Android vulnerability detection improvements

NEW INTEGRATION EJBCA certificate authority integration

NEW INTEGRATION Google Cloud Certificate Authority Service (CAS) integration

└──▷ BREAKING ON UPGRADE

- ! Fleet no longer sends Nudge settings to fleetd for hosts on macOS 13 and older; `fleetd`'s built-in Nudge integration is removed.
- ! `fleetctl gitops` and batch APIs now reject configuration profiles, scripts, reports, and labels whose names differ only by letter case, accents, or trailing spaces.
- ! `fleetctl gitops` on Fleet Free now returns an error for custom EST and Smallstep certificate authorities.
- ! Windows BitLocker key escrow now requires `server.private_key`.
- ! The Technician role can no longer run ad hoc scripts (`script_contents`) through the API or `fleetctl run-script`.
- ! The `FLEET_ALLOW_BOOTSTRAP_PACKAGE_DURING_MIGRATION` environment variable is deprecated in favor of `macos_enable_bootstrap_package_during_migration`.

WAS THIS USEFUL? 👍 👎

N-able N-central

Feedback ⓘ

N-central 2026.4 HF3 Release Notes (2026-11-16) BLOG ↗ RANK

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

> N-able N-central's 2026.4 HF3 release focuses on session-hardening: it adds a new Inactive Assets management tab and configurable appliance token behaviour for handling inactive Agents, while permanently blocking legacy APIs and requiring device identifiers for Agent connections.

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

01 **Inactive Assets tab for Agent management** Not checked yet NEW

70

A new Inactive Assets tab appears under Configuration > Asset Discovery > Discovered Assets, listing Agents flagged inactive after session-hardening changes (e.g. from cloned VMs or re-images), with every column filterable. Administrators can use Set Active, Delete, and Export (to CSV or PDF) on inactive Agents from this tab, with all actions logged for audit.

After upgrading Agents to HF3, review which Agents were marked inactive (e.g. from cloned VMs or re-images) and restore the correct one for a device.

📍 In N-central, go to Configuration › Asset Discovery › Discovered Assets, then select the Inactive Assets tab. Filter by device name, select the correct inactive Agent, and click Set Active to confirm.

— Names exact UI path and actions, lacks CLI/API equivalent

N-central 2026.4 HF3 Release Notes (202…

02 **Appliance token behaviour controls for inactive Agents**

Not checked yet NEW

60

New options in the Appliance token behaviour section of Administration > Mail and Network Settings > Network Setup let administrators control how N-central handles inactive Agents, such as requiring them to upgrade before reconnecting. Changes take effect immediately without a server restart.

Tune how N-central handles Agents it marks inactive — for example, requiring them to upgrade before reconnecting — without restarting the server.

📍 In N-central, go to Administration › Mail and Network Settings › Network Setup and locate the Appliance token behaviour section to configure inactive Agent

handling. Changes apply immediately.

— Exact settings path given, specific option names not enumerated

N-central 2026.4 HF3 Release Notes (202…

THINNER COVERAGE BELOW

03 Device identifier requirement for Agent connections

Not checked yet

BREAKING

50

N-central now rejects Agent connections when both the SMBIOS ID and machine ID are missing from a device, requiring at least one identifier to communicate with the server. Agents and Probes not upgraded to 2026.4.1 HF3 are susceptible to impersonation and are not fully protected by this session-hardening change.

— Clear mechanism and upgrade caveat, no remediation steps given

N-central 2026.4 HF3 Release Notes (202…

04 Legacy APIs permanently blocked

Not checked yet

BREAKING

30

The Legacy APIs configuration option has been removed and legacy APIs are now permanently blocked; any integration relying on legacy APIs will stop working after upgrading.

— States the change but no migration path or affected endpoint names

N-central 2026.4 HF3 Release Notes (202…

└──▷ **BREAKING ON UPGRADE**

! The configuration option for legacy APIs has been removed and legacy APIs are now permanently blocked — any integration relying on legacy APIs will stop working after upgrading.

! Devices without both a SMBIOS ID and a machine ID are now rejected by the N-central server — Agents on such devices will be unable to communicate until at least one identifier is present.

! Agents and Probes not upgraded to 2026.4.1 HF3 are susceptible to impersonation and are not fully protected by the new session-hardening changes.

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ

What's New in Iru: 13 New Auto Apps, Android Certificate Delegation & Outbound SCIM

NOTES ↗

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

Iru adds Android work profile certificate delegation and outbound SCIM to Iru Identity.

NEW SUPPORTED PLATFORM Iru now supports delegating certificate installation and selection permissions to managed apps in the Android work profile. Not checked yet

NEW INTEGRATION Iru Identity adds outbound SCIM support. Not checked yet

NEW INTEGRATION New Auto Apps: GitHub Copilot, Superhuman, Termius (Mac); Make Me Admin, Power BI Desktop, NetBird (Windows)

WAS THIS USEFUL? 👍 👎

◆ Attack Surface Management

Censys

Feedback ⓘ November 16, 2026 (2026-11-16) BLOG ↗

Censys is an internet intelligence platform that continuously scans and indexes internet-connected devices, services, and certificates to help security teams discover attack surfaces.

Censys adds `fields` and `exclude_fields` to asset lookup API endpoints, human-readable host reputation summaries, and updated Palo Alto Cortex and Google SecOps SOAR integrations.

Filter API responses on the Get host enrichment, Get host, Get web property, and Get certificate endpoints by passing `fields` to include only specific fields or `exclude_fields` to omit specific fields. Not checked yet

Generate human-readable summaries of host reputation scores to understand a host's label more easily. Not checked yet

Censys's Palo Alto Cortex integration (v3.2.4) now adds a Determine IP score using Censys Host Reputation data (for Censys Core customers) configuration parameter that scores IPs as Malicious, Suspicious, or Good based on the Censys host reputation label. Not checked yet

The Palo Alto Cortex integration now accepts custom labels in the IP Malicious labels and IP Suspicious labels parameters beyond the listed options, and also considers service labels alongside host labels when scoring. Not checked yet

The Google SecOps SOAR integration now supports enriching hosts using the get host enrichment API. Not checked yet

NEW DETECTION New software, hardware, and OS fingerprints added to the Platform

NEW DETECTION New ASM risk fingerprint: Certificate Issued by Compromised eMudhra Intermediate CA

WAS THIS USEFUL? 👍 👎

Qualys CyberSecurity Asset Management

Feedback

changelog · Nov 16, 2026 BLOG ↗

Qualys CyberSecurity Asset Management continuously discovers, inventories, and assesses all IT assets across hybrid environments to eliminate security blind spots and reduce attack surface risk.

Qualys ETM Identity adds an ISPM layer correlating identity exposure with Enterprise TruRisk, covering attack-path analysis, real-time AD monitoring, and closed-loop remediation.

NEW INTEGRATION NEW DETECTION Build a unified identity inventory across Active Directory, Entra ID, and Okta, covering human, machine, service, and shadow identities and their relationships. Not checked yet

Correlate identity risk with Enterprise TruRisk by combining Qualys Cloud Agent and native connector data with Identity TruRisk scoring. Not checked yet

Surface high-risk identity exposure including stale accounts, privileged identities, excessive access, and toxic combinations. Not checked yet

Understand attack paths and trust relationships through AI-summarized attack-path analysis and a Domain Trust Map showing how identity exposure can reach critical assets. Not checked yet

Monitor identity changes in real time with AD Real-Time Monitoring to catch changes that introduce new exposure. Not checked yet

+1 more

NEW INTEGRATION Remediation workflow integrations with Jira and ServiceNow

WAS THIS USEFUL? 👍 👎

◆ Patch Management

Automox

Feedback ⓘ

Manage Exclusion Windows Directly by Device (2026-11-16) BLOG ↗ RANK

Automox is a check out our full list of release notes, and stay up to date with what

> Automox replaced group-based exclusion window scheduling with direct per-device assignment and control, while deprecating the old group-scoped scheduling API in favor of new per-device endpoints.

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

01 Group-scoped exclusion window endpoints deprecated

Not checked yet BREAKING 85

The group-scoped scheduled window API endpoints — Create, Get, Update, Delete, List, Check Active, Check Status for Server Groups, List Next Dates by Group, and List Next Dates by Device — are deprecated and replaced by new per-device exclusion window API endpoints. Additionally, moving a device between groups no longer adds or removes its exclusion window, since exclusion windows are now managed independently of group membership.

Review the Automox API docs for deprecation notices on group-scoped endpoints and identify the new per-device exclusion window endpoints to migrate your automation scripts.

📍 steps

— Names all deprecated endpoints and migration path

Manage Exclusion Windows Directly by De…

THINNER COVERAGE BELOW

02 Per-device exclusion window management

Not checked yet NEW 58

Exclusion windows can now be assigned and managed directly on individual devices instead of through group assignments, using a new device-selection workflow in the Automox console. A new Edit Devices action lets admins add or remove devices from an exclusion window, and updated Device Details pages and tables now show which exclusion windows apply to each specific device.

— Names UI workflow and action but no config/API detail

Manage Exclusion Windows Directly by De…

└──▷ **BREAKING ON UPGRADE**

! Existing group-scoped scheduled window API endpoints (Create, Get, Update, Delete, List, Check Active, Check Status for Server Groups, List Next Dates by Group, List Next Dates by Device) are deprecated; migrate to the new per-device exclusion window endpoints.

! Moving a device between groups will no longer add or remove its exclusion window — exclusion windows are now managed independently of group membership.

WAS THIS USEFUL?  

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Ory

Feedback ⓘ changelog · Nov 17, 2026 CHANGELOG ↗ RANK

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

> Ory shipped granular outbound HTTP limits for Hydra, session risk rules for Enterprise projects, persistent SAML SP entity IDs, and custom OAuth 2.0 access token prefixes.

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

01 Per-feature outbound HTTP limits in Hydra

Not checked yet NEW

85

Ory Hydra v26.4.16 lets you configure outbound HTTP limits per feature using `default_http_client` for the global default and feature-specific `http_client` overrides (e.g. `max_conns_per_host`), so a single feature's outbound connections can be capped without affecting the global default.

Cap outbound HTTP connections for a specific Hydra feature without affecting the global default.

yaml

```yaml
default_http_client:
  max_conns_per_host: 100

some_feature:
  http_client:
    max_conns_per_host: 10
```

— Named config keys plus a runnable YAML example changelog-20261117-c30a9a3a

02 SAML SP entity ID storage

Not checked yet NEW

65

Each SAML provider on Ory Network now stores its SP entity ID in `sp_entity_id`, representing the audience its IdP sends, shipped in v26.4.14.

— Names exact field but no usage example or command changelog-20261117-c30a9a3a

THINNER COVERAGE BELOW

03 **Custom OAuth 2.0 access token prefix** Not checked yet NEW

45

Enterprise projects on Ory Network can now set a custom prefix for issued OAuth 2.0 access tokens on a per-project basis, added in v26.2.13.

— Describes capability but no config key or command shown changelog-20261117-c30a9a3a

04 **Session risk rules for Enterprise projects** Not checked yet NEW

30

Ory Network Enterprise plan projects can now configure session risk rules, introduced in v26.4.15.

— No mechanism or config surface given changelog-20261117-c30a9a3a

WAS THIS USEFUL?  

Clerk

Feedback ⓘ changelog · Nov 17, 2026 CHANGELOG ↗ 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's MCP server gained three new tools for listing and searching organizations, their memberships, and user-organization relationships, requiring reconnection to pick up a new permission scope.

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

01 Organization listing tools in MCP server

Not checked yet NEW

67

Three new MCP tools let agents query organization data: `list_organizations` lists or searches organizations with filters and pagination, `list_organization_memberships` lists an organization's members with their roles and permissions, and `list_user_organization_memberships` lists the organizations a user belongs to along with their role and permissions in each.

— Names all three MCP tools but no parameters or output schema

changelog-20261117-ac4b5997

02 New organizations:read permission required for MCP tools

Not checked yet

65

BREAKING

The new organization MCP tools require an `organizations:read` permission. Existing Clerk MCP server connections may be missing this permission — users should reconnect by signing out and connecting again, and can verify the permission is present by asking the agent to run `whoami`.

Verify your MCP connection has the required `organizations:read` permission before running any organization tools.

```
$ whoami
```

— Gives exact remediation step and verification command changelog-20261117-ac4b5997

└──▷ **BREAKING ON UPGRADE**

! Existing Clerk MCP server connections may be missing the new `organizations:read` permission required by the organization tools — reconnect by signing out and connecting again.

WAS THIS USEFUL?  

ConductorOne

Feedback ⓘ November 16, 2026 (2026-11-16) BLOG ↗

RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

> ConductorOne's biggest change this window is the c1i 0.10.1 CLI release, which adds new CRUD commands for entitlements, apps, and automations alongside a long list of breaking changes to output format, flags, and exit codes; the platform itself also gained AI-assisted connector sync debugging, no-expiry risk acceptance for findings, and OAuth2 token parameter support for custom MCP servers.

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

01 c1i output format and safety behavior changes

Not checked yet

98

BREAKING

Commands that create, update, or delete something now print the API response without its wrapper key, and a delete prints `{}` — e.g. `.app.id` becomes `.id` for `apps create` and `.policy.id` becomes `.id` for `policies create`. Commands that used to print a text line now print JSON. `--dry-run` now prints one JSON object `{"dryRun":true,"requests":[...]}` instead of text and shows credential values as `[redacted]`; setting `C1I_DRY_RUN=yes` or another unparseable value now causes a hard failure instead of silently sending the change. Passing `--fields` to a command that changes something is now an error. `--error-format json` now makes everything `c1i` writes to stderr NDJSON, including notices and progress, not just errors — filter with `jq 'select(.error)'`. `apps add-owner` and `apps remove-owner` now take the app as the first argument and the user as `--user-id`; the old argument order exits with an error and sends nothing. Invalid input (blank ID, ID of `.` or `..`, malformed input file, or an unparseable `C1I_DRY_RUN`, `C1I_DEBUG`, or `C1I_MAX_RETRIES` value) now exits `2`; stuck pagination exits `6`; failure reaching the docs site exits `8`.

Delete an app by ID and confirm the response is a bare object rather than a wrapped one, using the new 0.10.1 output format.

```
$ c1i apps delete <app-id> | jq '.id'
```

Preview a destructive automation change safely — credentials are redacted and nothing is sent.

```
$ c1i automations update <automation-id> --name 'New Name' --dry-run
```

Filter the error stream from a scripted bulk operation to surface only real failures.

```
$ c1i entitlements update <entitlement-id> --description 'Updated' --error-format json 2>&1 | jq 'select(.error)'
```

— Names every flag, exit code, and command change verbatim.

November 16, 2026 (2026-11-16)

THINNER COVERAGE BELOW

02 **Custom OAuth2 parameters for MCP server token requests**

Not checked yet NEW 55

Custom MCP servers using OAuth2 client credentials can now send extra token request parameters such as `audience` or `resource`, so providers like Auth0 no longer reject the token request.

— Names the parameters and provider but no config path. November 16, 2026 (2026-11-16)

03 **New c1i CRUD commands for entitlements, apps, and automations**

Not checked yet NEW 50

`c1i` v0.10.1 adds commands to read a single task and update, delete, or manage owners and webhooks for entitlements, apps, and automations, and to filter the users list by origin.

— Describes new commands but not their exact names. November 16, 2026 (2026-11-16)

04 **No-expiry option for risk-accepted findings** Not checked yet NEW 50

Findings can now be risk-accepted with no expiry by choosing No expiry under Accept for, keeping them accepted until manually reopened.

— Gives exact UI option but no further mechanism. November 16, 2026 (2026-11-16)

05 **AI-assisted connector sync debugging** Not checked yet NEW

50

A failing or count-dropping connector sync can now be debugged by asking the C1AI assistant via Debug sync with AI on the connector's details page.

— Names the UI entry point but not the debugging mechanism.

November 16, 2026 (2026-11-16)

06 **Richer shadow app details for Google Workspace** Not checked yet

IMPROVED

50

Shadow apps discovered through Google Workspace now show each account's last OAuth event, granted scopes, and last known IP address.

— Lists the new fields but no scope beyond Google Workspace.

November 16, 2026 (2026-11-16)

07 **Immediate entitlement linking for custom-created accounts**

Not checked yet IMPROVED

50

Accounts created with Custom values are now linked to their user immediately, so entitlement grants can follow right away without waiting for the next sync.

— Explains the before/after but no UI path given. November 16, 2026 (2026-11-16)

08 **Connector error detail in automation run info** Not checked yet

IMPROVED

50

When a connector action step fails inside an automation, the Run info drawer now shows the connector's own error message instead of only 'app action failed.'

— Names the exact drawer and old/new error text. November 16, 2026 (2026-11-16)

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! Commands that create, update, or delete something now print the API response without its wrapper key, and a delete prints `{}` — update any `jq` paths that read the wrapped id, such as `.app.id` to `.id` for `apps create` and `.policy.id` to `.id` for `policies create`.

! Commands that used to print a text line now print JSON.

! `--dry-run` now prints one JSON object `{"dryRun":true,"requests":[...]}` instead of text, and shows credential values as `[redacted]`.

! Passing `--fields` to a command that changes something is now an error.

! `--error-format json` now makes everything `c1i` writes to stderr NDJSON, including notices and progress, not just errors — select errors with `jq 'select(.error)'`.

! `apps add-owner` and `apps remove-owner` now take the app as the first argument and the user as `--user-id`; the old argument order exits with an error and sends nothing.

! Invalid input (blank ID, ID of `.` or `..`, malformed input file, or unparseable `C1I_DRY_RUN`, `C1I_DEBUG`, or `C1I_MAX_RETRIES` value) now exits `2`; stuck pagination exits `6`; failure reaching the docs site exits `8` — several other exit codes also changed.

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Embedded Authorize (Beta): in-app sign-in, with Auth0 powering each step (2026-11-12)

BLOG ↗ RANK

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

> Auth0 introduced Embedded Authorize BETA, a new API-driven way to build in-app sign-in flows without redirecting to Universal Login.

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

01 **Embedded Authorize beta for in-app sign-in** Not checked yet NEW

88

A new `POST /e/authorize` endpoint lets apps build in-app sign-in flows without redirecting to Universal Login; it returns the allowed next steps for the app to render until Auth0 issues an authorization code. During Beta it supports email and phone one-time code sign-in, and second-factor MFA via email OTP, phone OTP, TOTP, push, and recovery code. It has out-of-the-box compatibility with Auth0 Universal Components for web, iOS, and Android.

Kick off an in-app sign-in flow — no browser redirect required — and receive the first allowed step your UI should render.

```
$ curl -X POST https://<your-tenant>.auth0.com/e/authorize \
  -H 'Content-Type: application/json' \
  -d '{"client_id": "<client_id>", "redirect_uri": "
<redirect_uri>", "response_type": "code", "scope": "openid
profile email"}'
```

— Names endpoint, auth methods, MFA factors, and includes runnable example.

Embedded Authorize (Beta): in-app sign-…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Privileged Access Management

Teleport

Feedback ⓘ v18.11.10 NOTES ↗ RANK

Teleport is an open-source access platform that provides secure, identity-based access to SSH servers, Kubernetes clusters, databases, and web applications.

> Teleport v18.11.10 adds replay support for Beam-related audit events, including traffic captures and session end events.

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

01 **Replay support for Beam audit events** Not checked yet NEW

30

Teleport can now replay Beam-related audit events, including traffic captures and session end events, through the audit log.

— Names the event types but no mechanism or UI path given. v18.11.10

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Zama fhEVM

Feedback ⓘ v0.14.7 NOTES ↗ RANK

Fhevm is a zama uses Fully Homomorphic Encryption (FHE) so institutions can issue, trade and manage confidential assets on public blockchains, fully verifiable.

> fhEVM's v0.14.7 release adds access control to HCU limit enforcement and introduces relayer API key authentication to its end-to-end test suite.

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

01 **Allowlist for HCULimit host contracts** Not checked yet NEW

45

Adds an allowlist to HCULimit host contracts to control which addresses are subject to HCU limits.

— Names the contract but not allowlist mechanics or management. v0.14.7

02 **Relayer API key authentication in end-to-end tests**

Not checked yet NEW

25

Adds relayer API key authentication support to the end-to-end test suite.

— Bare statement, no key format, endpoint or config shown. v0.14.7

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v4.3.3 NOTES ↗ RANK

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 200+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

> CISO Assistant's v4.3.3 release adds a small workflow automation capability: inline event triggers that set field defaults automatically.

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

01 **Field defaults via inline workflow event triggers** Not checked yet

NEW 33

Workflows can now set field defaults through inline event triggers, reducing manual configuration when a workflow step fires.

— Describes mechanism but no exact config key or UI path given v4.3.3

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

cloudformation-guard

Feedback ⓘ 3.2.2 NOTES ↗ RANK

CloudFormation Guard is an open-source CLI tool that validates AWS CloudFormation templates against policy-as-code rules written in a custom DSL.

> CloudFormation Guard 3.2.2 adds two DSL conveniences — a `key` accessor and optional function naming — and splits the Guard language out into a standalone library separate from the CLI.

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

01 **New `key` method and function naming in Guard DSL** Not checked yet

NEW

55

Guard DSL rules gain a `key` method to access the key of a key-value pair during evaluation. Built-in Guard DSL functions can now also accept an optional function name parameter, intended to improve rule readability and error reporting.

— Names both DSL additions but no example usage shown. 3.2.2

02 **Guard DSL extracted as standalone library** Not checked yet NEW

25

The Guard language is now extracted as a standalone library, available separately from the CLI.

— Only a one-line mention, no package name or usage given. 3.2.2

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 6.0.0 (2026-11-16) BLOG ↗ RANK

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

> Runlayer 6.0.0 adds agent error notifications across Slack, Teams, and webhooks, starts logging refused OAuth token requests in audit trails, and tightens both the response format and validation of the policies API.

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

01 **Policy scope fields returned as lists** Not checked yet BREAKING

95

`GET`, `POST`, and `PUT` on `/api/v1/policies` now return `scope.servers`, `scope.tools`, and `scope.resources` as lists — `"*"` becomes `["*"]`. Clients that compare a scope field to `"*"` must also accept `["*"]`.

— Exact endpoint, methods, fields, and migration note given 6.0.0 (2026-11-16)

02 **Strict validation of policy scope keys** Not checked yet BREAKING

72

A request scope or agent account connection policy with an unknown key (e.g. a misspelled `tool`) now returns HTTP 422 instead of being silently ignored.

— Clear before/after behavior but endpoint not restated 6.0.0 (2026-11-16)

03 **Audit logs for refused OAuth token requests** Not checked yet NEW

69

Refused OAuth token requests now appear in audit logs with the client ID, agent or user involved, and the reason — covering wrong client secrets, failed PKCE checks, and replayed ID-JAGs.

— Specific failure causes named but no endpoint or log location 6.0.0 (2026-11-16)

04 **Agent error notifications to Slack, Teams, and webhooks**

Not checked yet NEW

67

Configure agent error notifications to Slack, Microsoft Teams, or signed HTTPS webhook destinations, with cooldowns and recent delivery history, for agent failures and unresolved tool errors.

— Names channels and cooldown mechanism but no exact config key 6.0.0 (2026-11-16)

└──▷ BREAKING ON UPGRADE

! `GET`, `POST`, and `PUT` on `/api/v1/policies` now return `scope.servers`, `scope.tools`, and `scope.resources` as lists — `"*"` becomes `["*"]`. Clients that compare a scope field to `"*"` must also accept `["*"]`.

! A request scope or agent account connection policy with an unknown key (e.g. a misspelled `tool`) now returns HTTP 422 instead of being silently ignored.

WAS THIS USEFUL?  

◆ Agent Sandboxes

Docker Sandboxes

Feedback ⓘ v0.54.0 NOTES ↗ RANK

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

> Docker Sandboxes v0.54.0 adds dedicated port management subcommands and a setting to disable post-quantum TLS key exchange for proxy-restricted networks, alongside several breaking security hardening changes affecting SSH config, clone-mode Git remotes, root home-directory resolution, and kit permission checks.

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

01 Port management subcommands for sandboxes

Not checked yet NEW

90

New commands let you manage published ports on a sandbox: list them with `sbx port ls <SANDBOX>`, publish a port with `sbx port publish <SANDBOX> <HOST_PORT>:<SANDBOX_PORT>` (omit the host port to assign one automatically), and remove mappings with `sbx port unpublish <SANDBOX> <SANDBOX_PORT>` or `<HOST_PORT>:<SANDBOX_PORT>` to remove a single mapping.

Publish a specific host-to-sandbox port mapping on a running sandbox, e.g. to expose a web server on port 8080.

```
$ sbx port publish my-sandbox 8080:8080
```

— Exact subcommand syntax with runnable example given v0.54.0

02 Disable post-quantum TLS key exchange

Not checked yet NEW

90

For firewalls or TLS-inspecting proxies that can't handle post-quantum key exchange, it can be disabled by setting `tls.postQuantumKeyExchange` to `false` via `sbx settings set tls.postQuantumKeyExchange false`, then restarting the daemon with `sbx daemon restart`.

Turn off post-quantum key exchange when a TLS-inspecting proxy is dropping HTTPS connections, then restart the daemon to apply the change.

```
$ sbx settings set tls.postQuantumKeyExchange false
sbx daemon restart
```

— Named config key and exact command sequence with example v0.54.0

03 **Clone-mode Git remote transport hardened** Not checked yet BREAKING

80

Clone-mode `sandbox-<name>` Git remote fetches now use the daemon socket instead of an unauthenticated localhost TCP port, requiring the sandbox to be running to fetch; the default transport also now rejects publishing sandbox TCP port `9418`. TCP compatibility for Git clients whose remotes ignore the `uploadpack` setting can be restored with `sbx settings set clone.legacyTcpRemote true`, which exposes the repository to other local accounts.

— Mechanism, risk, and exact opt-out setting all named v0.54.0

04 **SSH config regeneration required after upgrade** Not checked yet

65

BREAKING

Generated SSH configuration now prevents crafted host or user names from executing shell commands with older OpenSSH clients; existing configuration must be regenerated by rerunning `sbx setup ssh` (local) and `sbx --cloud setup ssh` (cloud) after upgrading.

— Security fix explained with exact remediation commands v0.54.0

05 **Root home directory resolution changed on Unix** Not checked yet

65

BREAKING

Running `sbx` as root on Unix now uses root's registered home directory; inherited overrides `HOME`, `XDG_CONFIG_HOME`, `XDG_STATE_HOME`, `XDG_CACHE_HOME`, `SANDBOXES_STORAGE_ROOT`, and `DOCKER_SANDBOXES_API` are ignored. Any older root daemon using custom storage must be stopped before upgrading.

— All affected env vars named plus upgrade precaution v0.54.0

THINNER COVERAGE BELOW

06 **Stable per-check IDs in diagnostics output** Not checked yet NEW

55

`sbx diagnose --format json` and `--format yaml` now include a stable `id` per check, suitable for matching checks in scripts.

— Command and formats named but no worked example v0.54.0

07 **MCP mode visibility in sandbox inspect** Not checked yet IMPROVED

50

`sbx inspect` now shows a local sandbox's static or dynamic MCP mode and its connected servers.

— Describes what's shown, clear command to check v0.54.0

---

08 **Clearer cloud API proxy error messages** Not checked yet IMPROVED

50

Cloud API connection failures behind HTTP proxies now report that the API is unreachable and suggest running `sbx --cloud diagnose`, instead of returning a generic 500 error.

— Before/after behavior and a follow-up command named v0.54.0

---

09 **Network policy check names matching rule** Not checked yet IMPROVED

45

`sbx policy check network` now names which network policy rule allowed a target, instead of just confirming access.

— Command named but behavior described briefly v0.54.0

---

10 **Working directory for scheduled sandbox runs** Not checked yet

NEW

45

Scheduled sandboxes can now run in a specific working directory using `--workdir` or `-w` on scheduled runs.

— Flag named but no example or further mechanism v0.54.0

---

11 **`sbx skills delete` alias** Not checked yet NEW

40

Adds the `sbx skills delete` alias for removing installed skills.

— Single named command with no further detail v0.54.0

---

12 **Architecture-aware cloud sandbox creation** Not checked yet

IMPROVED

40

Cloud sandbox creation now honors the requested architecture when combining published kits with locally built mixins.

— Behavior change stated without a command or example v0.54.0

---

13 **Kit permission re-confirmation from stronger digests**

Not checked yet

35

BREAKING

Existing kit permissions may require confirmation again after upgrading because permission checks now use stronger service digests.

— States impact but not the mechanism or digest type v0.54.0

└──▷ **BREAKING ON UPGRADE**

! Rerun `sbx setup ssh` (local) and `sbx --cloud setup ssh` (cloud) after upgrading — generated SSH configuration now prevents crafted host or user names from executing shell commands with older OpenSSH clients, and existing configuration must be regenerated.

! On Unix, running `sbx` as root now uses root's registered home directory; inherited overrides `HOME`, `XDG_CONFIG_HOME`, `XDG_STATE_HOME`, `XDG_CACHE_HOME`, `SANDBOXES_STORAGE_ROOT`, and `DOCKER_SANDBOXES_API` are ignored. Stop any older root daemon using custom storage before upgrading.

! Clone-mode `sandbox-<name>` Git remote fetches now use the daemon socket instead of an unauthenticated localhost TCP port; the sandbox must be running to fetch. Use `sbx settings set clone.legacyTcpRemote true` to restore TCP access (exposes the repository to other local accounts). Clone-mode sandboxes using the default transport now reject publishing sandbox TCP port 9418.

! Existing kit permissions may require confirmation again after upgrading because permission checks now use stronger service digests.

WAS THIS USEFUL?  

EVALUATE

◆ AI Observability & Evals

ongrid

Feedback ⓘ v0.21.0 NOTES ↗ RANK

> ongrid's AIOps engine gained traceable evidence in root-cause analysis reports, finer control over model reasoning and tool compatibility, the ability to approve policy-denied remediation commands, and support for persisting large tool results.

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

01 **Approval for policy-denied host commands** Not checked yet NEW 35

Operators can now approve policy-denied host commands, unblocking remediation actions that were previously rejected by policy.

— Clear capability but no UI path or command named v0.21.0

02 **Model reasoning-level and tool compatibility control**

Not checked yet NEW 33

The AIOps engine now supports model reasoning levels and tool compatibility selection, letting operators tune which model capabilities are used.

— Names a control but no flag or config key given v0.21.0

03 **Traceable evidence in RCA reports** Not checked yet NEW 28

The AIOps engine now adds traceable evidence to root-cause analysis reports, grounding each conclusion in observed data.

— Describes benefit but no mechanism or format named v0.21.0

04 **Persistence of large tool results** Not checked yet IMPROVED 28

The AIOps engine now persists and analyzes large tool results that previously exceeded in-memory limits.

— States behaviour change but no limit or storage detail v0.21.0

WAS THIS USEFUL? 👍 👎

GOVERN

◆ Model & Agent Protection

ToolHive

Feedback ⓘ v0.66.0 NOTES ↗ RANK

> ToolHive's v0.66.0 release hardens Docker workload isolation with an opt-in read-only root filesystem, strengthens ID-JAG token issuance with actor validation, and adds audit and telemetry refinements.

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

01 **ID-JAG minting from exchange grants with actor validation**

Not checked yet IMPROVED 45

ToolHive now mints ID-JAGs (identity JSON Assertion Grants) from exchange grants, names the acting agent inside the issued token, and validates the actor before issuance.

— Explains mechanism but no concrete API or command shown v0.66.0

02 **Read-only root filesystem for Docker workloads** Not checked yet

NEW 40

Workloads running on Docker can now opt in to a read-only root filesystem, hardening container isolation.

— Names the capability but no flag or config key given v0.66.0

03 **Audit event subjects for vMCP embedders** Not checked yet IMPROVED

30

vMCP embedders can now add audit event subjects to emitted audit records.

— States the capability only, no schema or API detail v0.66.0

04 **Server profile in anonymous usage metrics** Not checked yet

IMPROVED 15

Anonymous usage metrics now include the server profile.

— Bare statement with no further detail v0.66.0

WAS THIS USEFUL? 👍 👎

h5i

Feedback ⓘ v0.5.6 NOTES ↗ RANK

H5i is an Axum-based Rust web framework with Lean 4-verifiable application kernels that run on servers, bare-metal embedded devices, and GPUs, plus integrated security testing.

> h5i's v0.5.6 release adds three new CLI plugins for security testing: browser automation behind a policy-checking proxy, a websec workbench for replaying captured HTTP requests, and a recon ledger for crawling attack surfaces.

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

01 **Browser automation with policy-checking proxy** Not checked yet

NEW 90

Run browser automation sessions backed by a policy-checking proxy using `h5i browser open`, `h5i browser snapshot`, `h5i browser requests`, and `h5i browser audit` to inspect what the browser requested and what the proxy refused.

Audit a live target's HTTP traffic during a browser session, then inspect every request the browser made and what was refused.

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

— Four named subcommands plus a runnable multi-step example v0.5.6

02 **Recon ledger plugin for attack surface discovery** Not checked yet

NEW 85

Install the recon ledger plugin with `h5i plugin install recon --from ./h5i-recon` and discover attack surfaces by crawling a target with `h5i recon crawl --max-requests`, capping the number of requests made.

Crawl a target to map its attack surface, capping the run at 200 requests.

```
$ h5i plugin install recon --from ./h5i-recon
h5i recon crawl --max-requests 200
```

— Install and crawl commands with flag plus a runnable example v0.5.6

03 **Websec workbench plugin for replaying requests** Not checked yet

NEW 70

Install the websec workbench plugin with `h5i plugin install websec --from ./h5i-websec` to replay and modify captured HTTP requests via `h5i websec replay`.

— Exact install and replay commands given, no usage example v0.5.6

WAS THIS USEFUL? 👍 👎

Minimal Software Research minimal Feedback ⓘ

release-f7b6e364 NOTES ↗ RANK

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

> Minimal shipped a new `min secret` command for managing secrets, extended `min box` with listing, JSON/YAML output and a tamper-evident audit trail, and added configurable detach keys plus a daily release-check notice.

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

01 Box listing and show with JSON/YAML output Not checked yet NEW

75

List and show box types with `min box list` and `show` commands, supporting `-o json` or `-o yaml` output via `min box show -o json|yaml` — useful for exporting a box's full spec for CI ingestion.

Export the full spec for a box as JSON for downstream processing or CI ingestion.

```
$ min box show -o json my-box
```

— Exact flag and runnable command provided release-f7b6e364

02 Box audit trail with parent chain merge Not checked yet NEW

75

Inspect a box's tamper-evident audit trail and merge records from its parent chain using `min box audit --parent my-box`.

Inspect a box's tamper-evident audit trail and merge records from its parent chain.

```
$ min box audit --parent my-box
```

— Mechanism and exact command given, but limited elaboration release-f7b6e364

THINNER COVERAGE BELOW

03 Configurable detach keys for exec attach Not checked yet IMPROVED

25

Detach from an exec attach session using configurable detach keys via `min` attach.

— No key syntax or example shown release-f7b6e364

WAS THIS USEFUL?

◆ AI Security Posture & Discovery

Onyx Security

Feedback ⓘ announced · Nov 16, 2026 PRESS ↗

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

Onyx adds a Glean integration enabling agent discovery, automated governance reviews, intent-aware runtime policy enforcement, and browser-level Glean Assistant controls.

Connect Onyx to Glean via the Glean API for read-only discovery of agent inventory, ownership, permissions, tools, and data access. Not checked yet

Use Glean's runtime hooks to evaluate prompts, responses, and proposed tool executions against enterprise AI policies at supported checkpoints, blocking violations before execution. Not checked yet

Identify and block out-of-policy employee use of Glean Assistants in the browser through Onyx runtime controls. Not checked yet

NEW INTEGRATION Access the Onyx Secure AI Control Plane on the Glean Connector hub to manage Glean agents alongside third-party copilots, custom agents, and cloud AI services in a single inventory. Not checked yet

WAS THIS USEFUL?

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

WAF - Payload logging for WAF custom rules (2026-11-12) BLOG ↗ RANK

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

> Cloudflare WAF now supports payload logging for custom rules, letting Enterprise accounts capture matched request payloads under public-key encryption.

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

01 **Payload logging for WAF custom rules** Not checked yet NEW

67

Enterprise accounts can enable payload logging for WAF custom rules, encrypting matched payloads with an account-level public key or a zone-level public key, where the zone-level key takes precedence over the account-level one.

– Explains mechanism and key precedence but no exact UI/API steps.

WAF - Payload logging for WAF custom ru…

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

Arcade

Feedback ⓘ changelog · Nov 17, 2026 CHANGELOG ↗ RANK

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

> Arcade's latest release focuses on access control and integration hardening: project API keys gain scoping and deactivation, OAuth connections get named profiles, remote MCP servers get Client ID Metadata Document authorization, and five toolkit integrations (Google Drive, Google Docs, Gmail, Linear, GitHub) pick up new parameters and filtering — alongside breaking changes to deployment and the runner image.

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

01 Toolkit integration upgrades across Google, Linear, GitHub

Not checked yet IMPROVED 80

Google Drive's `query` parameter is now optional and searches every shared drive by default, flags incomplete results, and reports last editor and sharer. Google Docs' `EditDocument` can now return only named tabs instead of the whole document. Gmail's `GetThread` takes a flag to leave out each message's HTML body. Linear adds server-side filtering for `ListInitiatives` and `ListProjectComments` now returns discussion comments too. GitHub's `DownloadFile` on a directory now explains and lists contents, with GraphQL tools now working on GitHub Enterprise Server.

— Names exact API methods and parameters across five toolkits changelog-20261117-d5f9cc2f

02 Custom APT packages for deployed servers

Not checked yet NEW 80

Deployed servers can declare `apt_packages` in `.arcade/deploy.json` to install Debian packages at build time, so additional system packages are available on first pod start without recompiling.

Declare additional Debian packages your deployed server needs at build time, so they are available on first pod start without recompiling.

```json
{
  "apt_packages": ["libpq-dev", "libxml2"]
}
```

— Exact config key, file path, and runnable example given changelog-20261117-d5f9cc2f

03 **Pending deployment cap for self-hosted installs** Not checked yet

NEW

80

Self-hosted installs can cap pending deployments per project with `deployments.max_pending_per_project` in `engine.extraConfig`, preventing a CI pipeline or runaway automation from flooding a project's deployment queue.

Prevent a CI pipeline or runaway automation from flooding a project's deployment queue by capping how many deployments can be pending at once.

```yaml
deployments:
  max_pending_per_project: 5
```

— Exact config key and runnable example given changelog-20261117-d5f9cc2f

04 **Scoped and deactivatable project API keys** Not checked yet NEW

60

Admins creating a project API key now choose its scope — call tools, manage the project, or both — instead of always granting full access. Admins can also deactivate a project API key immediately without deleting it, and reactivate it without issuing a new secret.

— Describes scope options and lifecycle, no API endpoint named

changelog-20261117-d5f9cc2f

THINNER COVERAGE BELOW

05 **Nullable API key/gateway columns in tool execution history**

Not checked yet

55

BREAKING

The `tool-execution-history` API key/gateway columns are now nullable and required — existing integrations must handle null values.

— Names table and columns affected changelog-20261117-d5f9cc2f

06 **Client ID Metadata Document auth for remote MCP** Not checked yet

NEW

50

Authorization requests and responses now record an expected issuer, supporting Client ID Metadata Document authorization for remote MCP servers.

— Names mechanism but no config key or endpoint changelog-20261117-d5f9cc2f

07 **Registry tools for MCP catalog browsing** Not checked yet NEW

45

A gateway can expose two Registry tools letting a signed-in agent browse the full project MCP catalog, behind a per-gateway setting that is off by default.

– No tool names or UI path given changelog-20261117-d5f9cc2f

08 **Runner image drops legacy database drivers** Not checked yet

45

BREAKING

The deployed-server runner image's preinstalled unixODBC, libaio, GSSAPI, and Oracle Instant Client compatibility link are removed.

– Lists removed components, no replacement given changelog-20261117-d5f9cc2f

09 **New error kind for tool refusals** Not checked yet IMPROVED

40

Tool execution results now report a new error kind for deliberate tool refusals (bad input or business rules); the prior fatal-error kind is now reserved for unhandled crashes.

– No field or exact kind value named changelog-20261117-d5f9cc2f

10 **Named OAuth connection profiles** Not checked yet NEW 35

Organizations can turn on named connection profiles per OAuth provider so an end-user can switch accounts for the same toolkit.

– Thin description, no config surface named changelog-20261117-d5f9cc2f

11 **Account-based gateway binding removed** Not checked yet BREAKING

35

The old account-based gateway binding is dropped; any gateway still using it will be blocked until migrated.

– No specifics on migration path given changelog-20261117-d5f9cc2f

12 **Porter deployment provider removed** Not checked yet BREAKING

35

The `porter` deployment provider is removed.

– Names provider removed, no migration detail changelog-20261117-d5f9cc2f

13 **Gateway linking for remote MCP App resources** Not checked yet

NEW

25

A gateway can now link a remote MCP App's resources back to the gateway that served them.

– Single sentence, no mechanism detail changelog-20261117-d5f9cc2f

└──▷ BREAKING ON UPGRADE

- ! The `tool-execution-history` API key/gateway columns are now nullable and required — existing integrations must handle null values.
- ! The old account-based gateway binding is dropped; any gateway still using it will be blocked until migrated (see Upgrade notes).
- ! The `porter` deployment provider is removed (see Upgrade notes).
- ! The runner image's preinstalled unixODBC, libaio, GSSAPI, and Oracle Instant Client compatibility link are removed from the deployed-server runner image (see Upgrade notes).

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · November 17, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue