THE CYBER TOOLCHAIN

NO. 058

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 28, 2026 · EVERY DAY

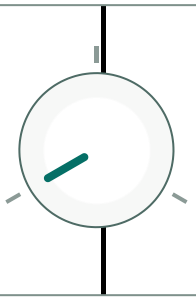
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 56 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

sslstrip

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

A tool for exploiting Moxie Marlinspike's SSL "stripping" attack.

> sslstrip 5.0.0 is a full Python 3.14 rewrite that replaces the old monolithic entry point with CLI subcommands, swaps iptables for a managed nftables firewall lifecycle, adds gated traffic logging and native gzip/Brotli/Zstandard decompression, and rebuilds the upstream networking stack on a shared async Twisted Agent with structured HTTP error handling.

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

01 **New CLI subcommand structure** BREAKING 75

sslstrip now exposes `sslstrip run`, `sslstrip doctor`, `sslstrip cleanup`, and `sslstrip --version` subcommands, replacing the old monolithic entry point; the top-level `sslstrip.py` script is removed entirely, so any code or scripts invoking it directly will break.

Run a pre-flight check to verify the environment is correctly configured before launching a strip session.

```
$ sslstrip doctor
```

— Names the four subcommands with a runnable example 5.0.0

02 **Gated traffic logging flags** NEW 70

Adds `--traffic-log` and `--traffic-log-file` flags for unredacted traffic logging; both flags are required together to enable logging, gating it from being turned on accidentally.

Capture unredacted HTTP traffic to a file during a strip session — useful for post-session credential review in a lab.

```
$ sslstrip run --traffic-log --traffic-log-file /tmp/strip-session.log
```

— Both flags named with a runnable command example 5.0.0

03 **nftables-based firewall lifecycle management** BREAKING

70

Adds managed nftables lifecycle with an exclusive `/run/sslstrip` lock, atomic JSON `nft` operations, privilege drop, and automatic forwarding restore, replacing the iptables-only setup; the iptables-based managed path is removed and nftables is now the first-class managed firewall backend.

— Mechanism and named lock path given, no direct command 5.0.0

04 **Structured HTTP error responses** NEW

70

Adds structured HTTP error responses: missing Host header → 400, CONNECT → 405, protocol upgrades → 426, upstream/DNS/TLS failures → 502, and timeouts → 504.

— Five status codes and triggers named precisely 5.0.0

05 **Shared async Twisted Agent for upstream requests** BREAKING

65

Adds a shared Twisted Agent with verified TLS policy and connection pool, orchestrated via native `async def` and `Deferred.fromCoroutine`; the old Twisted HTTPClient, factory, `DnsCache` client stack, and global singletons are removed.

— Mechanism and removed components named, no direct usage 5.0.0

06 **Lab CA file support for TLS verification** NEW

60

Adds a `--ca-file` flag to supply extra lab certificate authorities for TLS verification in controlled environments, used by the shared Twisted Agent 's verified TLS policy when connecting upstream.

Supply a lab CA bundle so the shared Twisted Agent can verify TLS to upstream servers that use a private PKI.

```
$ sslstrip run --ca-file /etc/lab/ca-bundle.pem
```

— Flag and command shown, no detail on CA validation logic 5.0.0

07 **Native gzip, Brotli, and Zstandard decompression** NEW

60

Adds native gzip, Brotli 1.2.0, and stdlib Zstandard decompression, including handling of stacked encodings, with an 8 MiB decoded-body limit.

— Formats, version and limit named, no usage path 5.0.0

08 **Python 3.14 minimum and locked packaging** BREAKING

60

Packaging now requires Python 3.14+ (`requires-python = ">=3.14"`) with exact pins for `Twisted[tls]` and Brotli via `uv_build` and a locked `uv.lock` ;

environments on Python 3.x below 3.14 will not work. mypy-based type checking is replaced by strict `ty`, breaking mypy-based CI or type-checking workflows.

— Exact version pins and config keys named 5.0.0

THINNER COVERAGE BELOW

09 **TTL and LRU bounds on secure-link store** IMPROVED 40

Adds TTL (30-minute) and LRU (10,000-entry) bounds on the secure-link store.

— Concrete numbers but no mechanism or usage detail 5.0.0

10 **Removal of legacy internals and shims** DEPRECATED 35

The dynamic import shim and camelCase modules are removed, breaking any code or scripts referencing them; favicon substitution and the hard-coded E*TRADE expression are also removed.

— Lists removed internals but no replacement guidance 5.0.0

└──▷ **BREAKING ON UPGRADE**

! Requires Python 3.14 or later (`requires-python = ">=3.14"`); any existing environment running Python 3.x < 3.14 will not work.

! The top-level `sslstrip.py` entry point is removed; invocation is now via the `sslstrip` CLI subcommands (`run`, `doctor`, `cleanup`).

! iptables-based managed path is removed; nftables is now the first-class managed firewall backend.

! Twisted `HTTPClient`, factory, `DnsCache` client stack, and global singletons are removed.

! Dynamic import shim and camelCase modules are removed; any code or scripts referencing them will break.

! Favicon substitution and the hard-coded E*TRADE expression are removed.

! mypy is replaced by strict `ty`; mypy-based CI or type-checking workflows will break.

WAS THIS USEFUL?  

Bishop Fox sliver

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

Adversary Emulation Framework

> Sliver v1.7.6 overhauls BOF and extension execution to run through Reflektor across Windows, macOS and Linux, while adding a c-archive implant output format, external-builder-driven profile generation, and configurable C2 keepalive — alongside a breaking change making multiplayer WireGuard opt-in.

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

01 **External builder generation from profiles** NEW 70

Profiles can now generate implants via external builders using `profiles generate external builder`, decoupling build configuration from the server.

— Exact command named, giving a clear starting point v1.7.6

02 **Reflektor-based execution for BOFs and extensions** IMPROVED 65

Default BOF execution now uses Reflektor instead of the previous in-process execution model, and implant extension loading also switched to Reflektor for reflective in-memory execution. Native extension support was added on macOS and Linux, previously Windows-only, broadening cross-platform BOF/extension use.

— Names mechanism and new platforms but no config flag or command shown v1.7.6

03 **c-archive output format for implants** NEW 60

Adds `c-archive` as a supported Go output format for implant generation, enabling implants to be embedded as C static libraries.

— Named output format but no build command example given v1.7.6

THINNER COVERAGE BELOW

04 **Non-interactive mcp start** IMPROVED 50

`mcp start` can now be run non-interactively, bypassing the interactive confirmation prompt.

— Exact command named though mechanism detail is minimal v1.7.6

05 **Multiplayer WireGuard now opt-in** BREAKING 45

Multiplayer WireGuard transport is now explicitly opt-in rather than enabled by default; existing setups relying on WireGuard being enabled by default in multiplayer mode will need to re-enable it manually.

— Breaking change flagged but no exact flag to re-enable named v1.7.6

06 **Configurable C2 keepalive interval** NEW 25

Adds configurable keepalive intervals for implant C2 communications.

— Thin bullet, no config key, flag or default value given v1.7.6

└──▷ **BREAKING ON UPGRADE**

! Multiplayer WireGuard is now explicitly opt-in; existing setups relying on WireGuard being enabled by default in multiplayer mode will need to re-enable it manually.

WAS THIS USEFUL? 👍 👎

◆ **Reconnaissance & OSINT**

webanalyze

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗ CODE ↗

Port of Wappalyzer (uncovers technologies used on websites) to automate mass scanning.

> webanalyze v0.4.5 adds a repeatable `-header` flag to inject custom HTTP request headers into fingerprinting scans.

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

01 **Custom HTTP headers via `-header` flag** NEW 80

A new repeatable `-header` flag (format `'Name: Value'`) sends custom HTTP request headers with every probe, enabling scans against hosts requiring auth tokens, custom user-agents, or other header-gated content.

Scan a host behind bearer-token authentication by injecting an Authorization header on every request.

```
$ webanalyze -host internal.example.com -crawl 2 -header
'Authorization: Bearer eyJhbGc...' -header 'X-Custom-Header:
value' -output json
```

— Names exact flag, format, and runnable example command. v0.4.5

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Keygraph shannon

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

Shannon is an AI pentester for web applications and APIs. It analyzes your source code, identifies attack vectors, and executes real exploits to prove vulnerabilities before they reach production.

> Shannon v2.7.0 adds support for running pentests using an xAI (Grok) subscription instead of pay-per-token API credits.

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

01 **xAI (Grok) subscription authentication** NEW 40

Shannon now supports authenticating pentests via an xAI (Grok) subscription plan, removing the need for pay-per-token API credits.

— Names the integration but no setup steps or config details. v2.7.0

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Veracode

Feedback ⓘ 1 RELEASE · 2026-08-28 CHANGELOG ↗

Veracode is a cloud-based application security platform that scans code for vulnerabilities and provides remediation guidance.

> Veracode's August 2026 update adds predependency scanning steps to the GitHub Workflow Integration, several findings and authentication improvements to the VS Code extension, and new language/framework coverage plus sharper credential-detection in the Static Analysis Engine.

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

01 **Scan findings filtering and OAuth region switching in VS Code**

IMPROVED 64

Veracode Scan for VS Code v1.16.3 adds a Current File filter to Static Scan Findings that scopes displayed findings to the file open in the editor, an option to filter and group Static Scan Findings by file name (with findings grouped by severity within each file), and an option to change the region during OAuth authentication, with the current or most recently used region shown on the OAuth screen.

— Three concrete UI features named but no exact navigation steps given

changelog-20260828-e9dd80d5

02 **Predependency commands in GitHub Workflow Integration** NEW

62

The Veracode YAML file for GitHub Workflow Integration v2.1.3 now supports custom `predependency` commands, enabling project preparation steps to run before a scan.

— Names config key and integration version but no example shown

changelog-20260828-e9dd80d5

THINNER COVERAGE BELOW

03 **Hardcoded credential detection improvements** IMPROVED

57

Static Analysis Engine v2026.08 adds detection of hardcoded values in `get` assignments modeled as getter functions, and improves detection of hardcoded passwords and credentials associated with `CWE-259` and `CWE-798`.

— Names specific CWE IDs and detection pattern but no user action described

changelog-20260828-e9dd80d5

04 **Language and framework support in Static Analysis Engine** NEW

Static Analysis Engine v2026.08 adds support for `Next.js 16.x`, `Android 16`, and the Python `google-cloud-bigquery` library.

— Names three concrete platforms/libraries but no scan mechanism detail

changelog-20260828-e9dd80d5

WAS THIS USEFUL? 👍 👎

OWASP dependency-track

Feedback

1 RELEASE · 2026-08-27

NOTES ↗

Dependency-Track is an intelligent Component Analysis platform that allows organizations to identify and reduce risk in the software supply chain.

> Dependency-Track 5.1.0 introduces initial KEV (Known Exploited Vulnerabilities) support across policies, notifications and metrics, adds new vulnerability data sources (Checkmarx SCA, JVN) and CycloneDX 1.7 ingestion, and improves the BOM upload API with an `isActive` flag, gzip-compressed uploads and workflow status polling.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 BOM upload API: isActive flag, gzip, workflow status NEW

86

Adds an `isActive` parameter to the BOM upload API endpoint, allowing projects to be marked active or inactive at upload time. Adds support for compressed SBOMs (e.g. gzip) in multipart BOM uploads, and exposes workflow run status on the BOM token polling endpoint.

Upload a compressed SBOM and mark the project active in one API call — useful for CI pipelines that produce gzipped CycloneDX BOMs.

```
$ curl -X POST 'https://dtrack.example.com/api/v1/bom' \
  -H 'X-Api-Key: <api-key>' \
  -F 'autoCreate=true' \
  -F 'isActive=true' \
  -F 'projectName=my-app' \
  -F 'projectVersion=1.0.0' \
  -F 'bom=@bom.cdx.json.gz;type=application/gzip'
```

— Named parameter plus runnable curl example 5.1.0

02 KEV (Known Exploited Vulnerabilities) support NEW

70

Implements initial support for CISA's Known Exploited Vulnerabilities (KEV) catalog, including use of KEV status in component and vulnerability policies and inclusion of a KEV flag in notifications. Adds REST API endpoints to view and trigger KEV mirroring runs, a KEV dimension to time series metrics, and a VulnCheck KEV data source.

— Policies, notifications and metrics named but no endpoint paths given 5.1.0

03 **DefectDojo v5 integration enhancements** IMPROVED 60

Forwards a `group_by` parameter to the DefectDojo (v5) integration and also forwards `analysis.detail` to DefectDojo (v5), enriching exported findings.

— Named fields but no full request example 5.1.0

THINNER COVERAGE BELOW

04 **New vulnerability data sources: Checkmarx SCA and JVN** NEW 40

Adds Checkmarx SCA as a vulnerability analyser and JVN (Japan Vulnerability Notes) as a vulnerability data source, expanding the upstream feeds Dependency-Track can ingest.

— Named sources but no configuration detail given 5.1.0

05 **GitHub App authentication for Advisories mirror** NEW 40

Adds GitHub App authentication as an option for the GitHub Advisories mirror.

— No configuration steps or keys given 5.1.0

06 **VEX import rating and semantics improvements** IMPROVED 35

Applies OWASP rating from VEX import, and aligns VEX semantics with CISA guidance and the broader tooling ecosystem.

— Prose description with no named endpoint or config 5.1.0

07 **Configurable dex and database query timeouts** IMPROVED 35

Makes the dex query timeout configurable and enforces a global database query timeout.

— No config key names or defaults given 5.1.0

08 **Red Hat support in distro-aware version matching** IMPROVED 35

Extends distro-aware version matching with Red Hat support, improving vulnerability matching accuracy for Red Hat-based distributions.

— No mechanism or matching detail beyond distro name 5.1.0

09 **Server version added to OpenAPI spec** IMPROVED 35

Adds the server release version to the OpenAPI spec info block.

— Names exact spec field but no example given 5.1.0

10 **CycloneDX 1.7 support** NEW 30

Adds support for CycloneDX 1.7 format for SBOM ingestion.

— Bare format name, no ingestion detail 5.1.0

11 **Protobuf publishing option for notifications** NEW 30

Adds protobuf as a publishing option for webhook notifications.

— Named format but no config example 5.1.0

12 **Component scope evaluation policy** NEW 25

Adds a new policy type to evaluate component scope.

— No mechanism or evaluation criteria described 5.1.0

13 **Portfolio findings query and scheduling performance** IMPROVED 25

Optimizes portfolio-wide findings queries and improves activity task scheduling performance for deep backlogs.

— No benchmark numbers or scope details given 5.1.0

14 **Package metadata resolution state tracking** NEW 20

Introduces state tracking for package metadata resolution.

— Bare description with no mechanism 5.1.0

15 **Configurable banner backend** NEW 15

Implements a banner backend for displaying configurable banners.

— No config keys or UI path given 5.1.0

WAS THIS USEFUL?  

cve-lite-cli

Feedback ⓘ 1 RELEASE · 2026-08-27 NOTES ↗

Fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

> cve-lite-cli added CI-friendly PR automation and diagnostics flags, plus one-click npm registry links in its HTML findings report.

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

01 **Automated PR creation and debug flags** NEW 85

Adds `--create-pr`, which after `--fix` commits changes and opens a GitHub pull request automatically (requires the `gh` CLI); `--base <branch>` sets the base branch for `--create-pr` and defaults to `main`; and `--debug` writes verbose runtime and network diagnostics to a timestamped log file for troubleshooting scan failures.

Automatically apply fixes and open a PR against a non-default branch in one step, useful in CI pipelines targeting a release branch.

```
$ cve-lite --fix --create-pr --base develop
```

Capture verbose network and runtime diagnostics when troubleshooting connectivity or unexpected scan failures.

```
$ cve-lite --debug
```

Automatically apply fixes and open a PR against a non-default base branch — useful in CI pipelines targeting a release branch.

```
$ cve-lite --fix --create-pr --base develop
```

Capture verbose network and runtime diagnostics to a timestamped log file when troubleshooting scan failures.

```
$ cve-lite --debug
```

— Exact flags, defaults, and requirement (gh CLI) with runnable commands. product docs

THINNER COVERAGE BELOW

02 **npm registry links in HTML report** IMPROVED 45

The HTML report's findings table now links each package name directly to its npm registry page, enabling one-click research from scan output.

— Describes the change but no exact navigation path or command given. v1.31.0

WAS THIS USEFUL? 👍 👎

HashiCorp terraform

Feedback ⓘ 1 RELEASE · 2026-08-26

NOTES ↗

Terraform enables you to safely and predictably create, change, and improve infrastructure. It is a source-available tool that codifies APIs into declarative configuration files that can be shared amongst team members, treated as code, edited, reviewed, and versioned.

> Terraform v1.16 adds JSON output to more CLI commands, a `store` block for ephemeral values in `terraform_data`, expanded resource lifecycle and Actions trigger controls, module-level `import` blocks, and a provisioner bug fix that changes how `bastion_host_key` is applied.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 Prevent resource destruction via lifecycle block NEW

75

Adds `destroy = false` to resource `lifecycle` blocks to prevent a resource from ever being destroyed, e.g. `lifecycle { destroy = false }` inside a resource block.

Prevent a critical resource from accidental destruction by setting `destroy = false` in its lifecycle block.

hcl

```hcl
resource "aws_s3_bucket" "critical" {
  bucket = "my-critical-bucket"

  lifecycle {
    destroy = false
  }
}
```

— Exact runnable HCL snippet included v1.16.0

02 Action trigger failure handling and lifecycle hooks NEW

70

Adds `on_failure` modes `halt`, `taint`, and `continue` to resource action triggers, adds `before_destroy` and `after_destroy` event support to Actions, adds a `caller` symbol in Action configurations exposing the calling resource's object value, and allows `-invoke` to be combined with `-target` to specify the calling resource instance when multiple resources trigger the same action.

— Names all flags/keywords but no runnable example v1.16.0

03 **Mermaid format for dependency graphs** NEW 65

Adds `-format=mermaid` flag to `terraform graph` to output the dependency graph in Mermaid format, e.g. `terraform graph -format=mermaid > graph.md` for embedding in docs or wikis.

Visualize the Terraform dependency graph in Mermaid format for embedding in docs or wikis.

```
$ terraform graph -format=mermaid > graph.md
```

— Exact runnable command included v1.16.0

04 **Automatic hostname inference for terraform stacks** IMPROVED 60

The `terraform stacks` command now automatically infers the target hostname from `credentials.tfrc.json` when neither `TF_STACKS_HOSTNAME` nor `TF_CLOUD_HOSTNAME` is set.

— Names env vars and file, no runnable command v1.16.0

THINNER COVERAGE BELOW

05 **JSON output for state and workspace commands** NEW 55

Adds `-json` flag to `terraform state show` and `-json` flag to `terraform workspace list` for machine-readable output, e.g. `terraform state show -json aws_instance.web`.

Get machine-readable JSON output of a resource's state for scripting or audit pipelines.

```
$ terraform state show -json aws_instance.web
```

— Two named flags, one runnable example v1.16.0

06 **Module scoping in terraform console** NEW 55

Adds `-scope=<module address>` flag to `terraform console` to evaluate expressions within a specific module or module instance scope.

— Names flag but no worked example v1.16.0

07 **Ephemeral value storage in terraform_data** NEW 55

Adds a `store` block to the `terraform_data` resource to hold ephemeral and sensitive values across plan and apply.

— Names block and behavior, no example given v1.16.0

08 **Fixed bastion_host_key application in provisioners** BREAKING

55

`bastion_host_key` is now correctly applied by provisioners, correcting prior incorrect behavior. Users should review their provisioner configurations to verify the configured key is correct before upgrading.

— Names the field and migration caution, no exact steps v1.16.0

09 **Module-level import blocks** NEW

45

Supports `import` blocks inside modules, rather than only at the root module level.

— Names the surface, thin on mechanism v1.16.0

10 **Policy evaluation summary for plan and apply** NEW

40

Terraform now displays a summary of policy evaluation outcomes for `plan` and `apply` runs against HCP Terraform.

— Describes what appears, no navigation detail v1.16.0

11 **Linux s390x binary support** NEW

40

Terraform is now available as a pre-built binary for Linux s390x (zLinux).

— Clear download target, no install command shown v1.16.0

12 **Resource reporting for skip_cleanup in tests** IMPROVED

40

Terraform now reports which resources were left behind when `skip_cleanup` is set in tests.

— Names config key, no example output shown v1.16.0

13 **Null support in contains() function** IMPROVED

35

The contains() function can now test for `null` values.

— Names function, no example usage v1.16.0

14 **Provider planned private data persistence** IMPROVED

30

Terraform now stores planned private data for providers, preserving provider-specific state across plan and apply.

— Internal behavior change, no user-facing surface named v1.16.0

15 **Nested blocks as computed provider values** NEW

25

Providers can now use nested blocks as computed values.

— Bare description, no mechanism or example v1.16.0

└──▷ **BREAKING ON UPGRADE**

! `bastion_host_key` is now correctly applied by provisioners — review your provisioner configurations to verify the configured key is correct before upgrading.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ 1 RELEASE · 2026-08-28 BLOG ↗

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

> GitHub is merging Copilot Chat, cloud agent, and mobile into a single unified experience and policy, while also shifting the default code review effort level and moving Copilot Business/Enterprise billing to an upfront model.

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

01 Code review default effort level shifts from Lite to Balanced

90

BREAKING

Starting September 28th, 2026, the default Copilot code review effort level changes from `Lite` to `Balanced` for all repositories and organizations; repos or orgs relying on `Default` meaning `Lite` must explicitly set `Lite` before that date. Organization-level and repository-level defaults work independently — the org default applies to repos without their own setting, while the repo default governs automatically requested reviews — and for manually requested reviews the effort level (`Lite`, `Balanced`, etc.) can be chosen directly from the `Reviewers` bar in the pull request UI.

Keep `Lite` as your org's code review default before the September 28th cutover, so automated reviews don't silently become more expensive.

> 📍 Go to your organization's Copilot settings on github.com › Code review › Review effort level, then change the value from `Default` to explicitly select `Lite` before September 28th, 2026.

— Names exact settings, UI path, and effort-level values with clear before/after.

Upcoming changes to GitHub Copilot poli…

02 Unified Copilot experience across Chat, Mobile, cloud agent

75

BREAKING

Copilot Chat on github.com, Copilot Chat in GitHub Mobile, and Copilot cloud agent are merging into a single unified experience and policy, enabled by default no earlier than September 28th, 2026. As part of this shift, chat data retention extends from 28 days to the life of the account under the agent sessions model, and the cloud agent gains a Sandbox for a faster cloud execution environment. Opting out of the unified experience after September 28th, 2026 removes access to Copilot on github.com and GitHub Mobile entirely, and organizations must confirm the policy via Copilot settings > `Copilot cloud agent (coming soon)` before that date.

Confirm or update the unified Copilot experience policy for your organization before it launches, to control whether members retain access on github.com and Mobile.

> 📍 Go to Copilot settings on github.com › Select `Copilot cloud agent (coming soon)` › Set the policy option to your desired value before September 28th, 2026.

— Names the merged surfaces and dates but mechanism of the merge itself is thin.

Upcoming changes to GitHub Copilot poli…

THINNER COVERAGE BELOW

03 Copilot Business/Enterprise billing moves to upfront per-seat charges

50

BREAKING

Copilot Business and Copilot Enterprise sign-ups reopen September 1st, 2026 with upfront per-seat billing required before users gain access; existing customers are moved to this model on October 1st, 2026, when all existing seats begin incurring an upfront charge at the start of each billing cycle rather than being billed in arrears.

— Gives exact dates and billing change but no configuration steps.

Upcoming changes to GitHub Copilot poli…

└──▷ BREAKING ON UPGRADE

- ! Starting September 28th, 2026, the `Default` review effort level for Copilot code review maps to `Balanced` instead of `Lite`; any repo or org relying on `Default` to mean `Lite` must explicitly select `Lite` in settings before that date or reviews will run at higher effort.
- ! Opting out of the unified Copilot experience after September 28th, 2026 removes access to Copilot on github.com and GitHub Mobile entirely.
- ! Starting October 1st, 2026, all existing Copilot Business and Copilot Enterprise seats incur an upfront charge at the start of each billing cycle rather than being billed in arrears.

WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ 2 RELEASES · 2026-08-28 BLOG ↗

JFrog is a DevOps platform providing software supply chain management, artifact management, and continuous integration/delivery automation tools.

> JFrog introduced a Preview early-access track, letting customers opt into new capabilities before general availability directly from the MyJFrog portal.

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

01 **JFrog Preview early-access track** NEW 72

A new opt-in early-access track offers two tiers: Closed Preview (invite-only feedback cohorts) and Open Preview (available to any eligible customer, fully supported), both free. Activate it in the MyJFrog portal under `Settings > Preview` by enabling the Enable Preview toggle, accepting the Preview Terms of Service in the Enable JFrog Preview dialog, and then selecting specific preview features to activate on your subscription.

Activate the JFrog Preview track to gain immediate access to newly available capabilities on your subscription.

> 📍 In MyJFrog, go to Settings › Preview. Click the Enable Preview toggle. In the Enable JFrog Preview dialog box, open and read the Preview Terms of Service. Select the checkbox to accept the terms and conditions, and click Agree and Enable. Then select the specific preview features you want to activate.

Opt in to JFrog Preview to gain immediate access to newly released capabilities before they reach general availability.

> 📍 In MyJFrog, go to Settings › Preview. Click the Enable Preview toggle. In the Enable JFrog Preview dialog box, open and read the Preview Terms of Service. Select the checkbox to accept the terms and conditions, and click Agree and Enable. Then select specific preview features to activate in your subscription.

– Names exact UI path and toggle but no specific preview features listed.

Introducing JFrog Preview: Getting New … launch-20260828-b987b724

WAS THIS USEFUL? 👍 👎

Socket

Feedback

3 RELEASES · 2026-08-25 → 2026-08-28

BLOG ↗

Socket detects supply chain attacks by analyzing open source package dependencies for malicious code, vulnerabilities, and suspicious behaviors.

> Socket extended its supply chain monitoring to Microsoft Edge extensions and launched two new ticketing integrations — ClickUp and Asana — that turn Socket alerts into bi-directionally synced tasks.

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

01 **Socket for ClickUp alert-to-task integration** NEW 90

New ClickUp integration (beta) lets teams create ClickUp tasks from Socket alerts via `Settings → Integrations → ClickUp`. Automated ticketing rules route matching findings into a chosen ClickUp Workspace and List with configurable priority (Urgent , High , Normal , Low , or Auto to map Socket priority automatically), tags, and assignee. Bi-directional status sync keeps alerts and tasks aligned — clearing an alert can move the linked task to completed, and closing or completing a task in ClickUp can triage-ignore the corresponding Socket alert — with updates delivered via Workspace webhooks. Manual task creation is available per alert via the Create Ticket button, selecting a target Workspace and List shown as `Space / List` or `Space / Folder / List`.

Set up the ClickUp integration in Socket to start routing alerts as tasks into your team's existing ClickUp workflows.

> 📍 In the Socket dashboard, go to Settings → Integrations → ClickUp and click Connect.

— Exact UI path and named config options given launch-20260826-99908cdf

02 **Socket for Asana alert-to-task integration** NEW 88

New Settings → Integrations → Asana UI lets an Owner or Admin connect one or more Asana workspaces via OAuth with no separate app to install. A Create Ticket button on any alert manually creates a linked Asana task, selecting workspace, project, tags, and assignee, with a bidirectional link between alert and task. Automated ticketing rules let teams select triggering alert events, add conditions (e.g. type = supply chain, priority = high), choose the target workspace and project, set tags/assignee, and configure two-way sync behavior. Two-way sync marks the Asana task complete when an alert clears, and completing the task can triage-ignore the linked Socket alert.

Connect an Asana workspace to your Socket organization so alerts can be routed as trackable tasks.

📍 In the Socket dashboard, go to Settings → Integrations → Asana and click Connect. Authorize via OAuth — no third-party admin console required. The connected workspace name, connecting user, and status will appear in the Socket dashboard.

Set up an automated ticketing rule to route high-priority supply chain alerts directly to the right Asana project without manual triage.

📍 In the Socket dashboard, go to Settings → Integrations → Asana → ticketing rules. Step 1: select the alert events that trigger the rule. Step 2: add conditions (e.g., type = supply chain, priority = high). Step 3: select the connected Asana workspace and target project. Step 4: add tags and assign to an Asana user. Step 5: configure two-way sync behavior — e.g., mark Asana task complete when alert clears, and triage-ignore the Socket alert when the Asana task is completed.

— Exact UI path and step-by-step rule setup given `launch-20260825-ff8c843b`

---

03 Extension monitoring for Microsoft Edge Add-ons NEW

70

Adds continuous analysis of extensions published through Microsoft Edge Add-ons, covering malware, credential and clipboard theft, data exfiltration, remote code loading, obfuscation, and impersonation — the same detections previously available for Chrome and Firefox. Enables version-change comparison, surfacing meaningful differences in code, permissions, network activity, and behavior between releases under the same extension identity. Identifies command-and-control infrastructure, suspicious endpoints, remote payload sources, and changes in contacted domains, and connects related extensions across stores through reused code, shared infrastructure, publisher patterns, and coordinated campaigns. This pairs with existing Chrome Web Store analysis to close the cross-store coverage gap.

Enable Edge extension coverage for your enterprise org — reach out to your Socket account team or contact Socket to get started with the Experimental rollout.

📍 In the Socket dashboard, contact your Socket account team to enable Edge coverage under the Experimental tier for your enterprise organization.

— Rollout is via account team request, not self-serve config `launch-20260828-da7c1bb0`

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL?

SonarSource SonarQube

Feedback

1 RELEASE · 2026-08-27

BLOG ↗

Trusted by 7M+ developers, SonarQube is the AI code review and code verification platform that catches quality and security issues before they reach production.

> SonarQube Cloud adds Hunter Agent, an AI security agent that detects broken access control, business logic, and authentication/session flaws that traditional SAST cannot evaluate, surfacing validated findings directly in the existing SonarQube triage workflow.

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

01 **Hunter Agent for logical flaw detection** NEW 79

Hunter Agent is now GA on SonarQube Cloud (Enterprise plan), an AI security agent that detects broken access control (IDOR, missing authorization checks, privilege escalation, CSRF), business logic flaws (skipped workflow steps, missing rate limits), and authentication/session management flaws (session fixation, non-expiring sessions, weak password recovery, missing MFA) — vulnerability classes SAST structurally cannot evaluate. Every candidate finding is validated for exploitability via a dedicated subagent before surfacing, achieving 80–90% average precision, and confirmed findings appear directly as SonarQube issues in the existing triage workflow with no separate portal, dashboard, or installation required. It supports on-demand triggering or scheduled runs (weekly or monthly) against a main branch, delivering findings asynchronously once confirmed; scales to codebases beyond 10 million lines of code via a structured multi-phase pipeline (Analyze, Explore, Validate, Synthesize) that cuts re-scan time on large repositories by roughly half through incremental scanning; works across the 40+ programming languages already supported by SonarQube with no new configuration beyond pointing it at an existing SonarQube project; and every finding carries an auditable discovery path showing how the agent reached its conclusion.

— Rich mechanism and named flaw classes, but no exact command or config key given.

launch-20260827-9c110bbd

WAS THIS USEFUL?

StepSecurity

Feedback ⓘ 2 RELEASES · 2026-08-27 BLOG ↗

StepSecurity is a security platform that hardens CI/CD pipelines by reducing attack surface and detecting suspicious activity in GitHub Actions workflows.

> StepSecurity's Dev Machine Guard agent (v1.16.0) adds two new fleet-wide inventory views — one surfacing plaintext credential exposure and one surfacing browser extension risk — to speed targeted incident response on developer machines.

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

01 Credentials inventory page in Dev Machine Guard NEW

87

A new Credentials page in the StepSecurity dashboard aggregates credential source coverage fleet-wide, showing per-source device count, plaintext vs. protected breakdown, location count, and entry count, filterable by category (Cloud, Source control, Package registries, Containers, Infrastructure) and exportable as CSV. It covers thirteen credential source types across five categories — `~/.npmrc`, `~/.pypirc`, `~/.vault-token`, `~/.netrc`, `~/.aws/credentials`, `~/.aws/config`, `~/.docker/config.json`, gcloud application default credentials, kubeconfig, `~/.ssh` private key files (`id_rsa`, `id_ed25519`, `id_ecdsa`), and `~/.git-credentials` — with a per-source detail view listing affected devices labelled Plaintext, Protected, or both and direct links to each machine. On macOS, Full Disk Access must be granted to the agent via MDM to reach TCC-protected credential paths for complete results. Available from agent `v1.16.0`; existing deployments gain the page automatically after each device's next scan.

When a credential-stealing package is disclosed, confirm which sources existed on affected machines so rotation is targeted rather than fleet-wide.

📍 In the StepSecurity dashboard, go to Dev Machine Guard › Credentials, filter by the relevant category (e.g. Cloud or Package registries), identify affected devices and entry counts, then export the view as CSV to drive a targeted rotation plan.

— Names exact credential paths, categories, macOS caveat and agent version.

launch-20260827-052c1615

02 Browser Extensions inventory page in Dev Machine Guard NEW

87

A new Browser Extensions page in the StepSecurity dashboard shows a fleet-wide table of every detected extension — name, extension ID, browser, store, version, device

count, and last-seen timestamp — filterable by browser, capability, store, and engine, with an Export CSV download. Raw browser permissions are rolled into ten named capabilities (`Global website access` , `Global request observation` , `Script injection` , `Global script injection` , `Cookie access` , `Global cookie access` , `Browsing history` , `Debugger access` , `Clipboard read` , `Extension management`) reflecting what the browser currently honours rather than just what the extension's manifest declares. A per-device detail panel lists affected devices grouped by version (first 500 devices shown) for rapid scoping against version-specific advisories. Covers Chrome, Edge, and Firefox on macOS, Windows, and Linux — including Firefox under snap and Flatpak and Edge under Flatpak — available from agent `v1.16.0` with no changes required for existing deployments.

During an incident, search by extension ID in the Browser Extensions page to instantly see which devices are running a compromised extension and at which versions.

> 📍 In the StepSecurity dashboard, go to Dev Machine Guard › Browser Extensions › search the extension ID in the table › review the per-device version grouping for affected machines.

— Names the ten capability labels, browsers, OS coverage and agent version.

`launch-20260827-a61ff06e`

WAS THIS USEFUL? 👍 👎

Aikido Security

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

Aikido Security is a developer-focused security platform that scans code, infrastructure, and dependencies for vulnerabilities and misconfigurations.

> Aikido launched agentic pentesting for Android apps, letting AI agents install a real APK on a test device, log in, and attack it over ADB with exploit-confirmed findings, joint app+backend coverage, and an AutoFix-and-retest workflow.

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

01 **Agentic Android (APK) pentesting** NEW 81

Aikido's agents install a real APK on a test device, log in as a user, and attack the app over ADB, with every finding exploit-confirmed before it appears in the report and reproduction steps including the exact ADB commands needed to replay the exploit. Coverage includes unsafe WebView usage, deep-link and exported-component abuse, client-side authentication issues, injection flaws, broken access control, weak session handling, unsafe API behavior, IDORs, cross-tenant access, and authentication bypasses. Assessments are rightsized by Aikido's indexer to the actual app size rather than charged a flat fee, and each run produces a recon report and attack plan with diagrams, threat scenarios, and the full agent reasoning trace (including vulnerabilities attempted but not confirmed). The capability is currently white-box only, requiring source repository access alongside the APK (black-box support is coming soon), and the uploaded build must have certificate pinning, root detection, emulator detection, and RASP disabled.

.
.
.
.

— Detailed mechanism and named requirements, but no runnable command for the reader

launch-20260826-3c11bdb8

02 **AutoFix PR and retest workflow for Android findings** NEW 73

After merging an AutoFix pull request for a confirmed Android finding, users can retest against the patched build without rescoping: open the original Android assessment in the Aikido console, go to the overview tab, upload the rebuilt APK, then click Retest on any finding from the initial pentest.

After merging an AutoFix PR for a confirmed Android finding, retest against the patched build without rescoping.

> 📍 In the Aikido console, open the original Android assessment, go to the overview tab, upload the rebuilt APK, then click Retest on any finding from the

initial pentest.

— Exact UI navigation and steps a reader can follow immediately launch-20260826-3c11bdb8

THINNER COVERAGE BELOW

03 **Joint app and backend API pentesting in one run** NEW

33

Aikido supports simultaneous assessment of the Android app and its backend API in a single run, with API findings appearing in the same report as app findings.

— States the capability but no mechanism or scope detail beyond one sentence launch-20260826-3c11bdb8

WAS THIS USEFUL? 👍 👎

Manifest

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

Automate SBOM generation, secure your software supply chain, and gain complete technology transparency across code, AI, vendors, and embedded systems." name="description"/> <meta content=

> Manifest launches Supplier Risk, a supplier-centric view of vendor risk grounded in SBOM data, adding a CLI flag to tag uploads by supplier, binary-generated SBOMs for vendors who can't provide one, and new portfolio, detail, and scoring views for third-party risk.

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

01 **Supplier flag for CLI uploads** NEW 75

Adds a `--supplier` flag to the Manifest CLI for associating an uploaded SBOM or binary directly with a named supplier at publish time.

— Names the exact CLI flag and its effect. launch-20260826-b483df8e

02 **Supplier Risk portfolio, detail, and scoring views** NEW 67

New supplier portfolio view ranks every third-party vendor by risk score, asset count, open vulnerabilities, and last-updated date in a single screen. A companion supplier detail view surfaces a vendor's full footprint — every asset, version, risk score, component count, and vulnerability count — in one click. Aggregated supplier risk scoring rolls up asset-level risk into a single per-supplier score for direct vendor-to-vendor comparison, and prioritized supplier findings stack-rank vulnerabilities by severity and product impact across a supplier's portfolio. The vulnerability and component pages were also extended to surface third-party products alongside first-party assets, with clear visual distinction between what teams build and what they buy.

— Concrete metrics and views named, but only UI navigation, no endpoint.

launch-20260826-b483df8e

03 **Binary-generated SBOMs for suppliers** NEW 62

Introduces binary-generated SBOMs for vendors who cannot provide one, covering C++, unmanaged code, and embedded systems where traditional SCA falls short, with results folded into that supplier's record.

— Describes scope and mechanism but no exact command given. launch-20260826-b483df8e

WAS THIS USEFUL? 👍 👎

sbomify

Feedback

2 RELEASES · 2026-08-25

BLOG ↗

sbomify is a product security artifact hub and a trust center.

> sbomify's product added Security Advisories with a gated Trust Center, a slate of CRA compliance tooling (patch SLA KPIs, risk_accepted state, EU Authorized Representative capture, export retention floors, EOL readiness checks), four new CSV exports, and a redesigned Trust Center UI, while sbomify-action cut its container image from 515MB to 106MB via on-demand digest-pinned tool fetching and added a public ClearlyDefined caching proxy plus support for eight more ecosystems.

└──▷ **WHAT SHIPPED · 20 FEATURES** most completely described first ? what's the number?

01 **clearly-cached ClearlyDefined caching proxy** NEW 93

Introduces `clearly-cached`, a public caching proxy at `clearly-cached.sbomify.com` that sits in front of ClearlyDefined and returns a flat 0.4KB projection of license fields, resolving 76 of 76 packages in 5s warm vs 31 of 76 in 474s direct. `SBOMIFY_CLEARLY_CACHED_URL` env var points the action at a self-hosted or local instance instead of the default.

Point sbomify-action at a self-hosted clearly-cached instance to avoid a public-internet dependency during enrichment.

yaml

```yaml
- uses: sbomify/sbomify-action@v26.8.0
  env:
    SBOMIFY_TOKEN: ${{ secrets.SBOMIFY_TOKEN }}
    SBOMIFY_CLEARLY_CACHED_URL: 'http://clearly-
cached.internal:8080'
```

— Named env var, endpoint, and concrete before/after benchmark numbers

launch-20260825-ae80c5f8

02 **Security Advisories module with visibility controls** NEW 90

New Security Advisories module lets you author vendor advisories with CVSS vector, affected/fixed version ranges, product associations, and a timestamped timeline, published directly to the Trust Center. Advisories support three visibility states — Private, Public, and Gated — where Gated requires both a workspace-level NDA grant and a link to a product the reader can already access; per-advisory product filtering redacts names the reader is not entitled to see.

Publish a gated security advisory so only customers under NDA for a specific product can read it, without revealing the advisory's existence to others.

> 📍 In the Security Advisories module, create a new advisory with severity, CVSS vector, affected version range (e.g. `>= 11.2.0, < 11.2.4-h20`), fixed version range, and timeline entries. Set visibility to `Gated`, attach it to the relevant product, then publish to the Trust Center. Readers without a workspace-level NDA grant and access to that product receive a 404.

— Full field list and gating mechanism, plus a walkthrough example

launch-20260825-b3c23b90

---

03 **Slimmed action image with on-demand tool fetching** IMPROVED

88

Adds `SBOMIFY_FETCH_RUNTIMES=0` env var to opt out of on-demand tool fetching for air-gapped environments; per-ecosystem tool bundles are now fetched from `sbomify/sbom-tools` at runtime, pinned by digest, verified, and attested, replacing tools baked into the image, dropping the container image from 515MB to 106MB. The generator fallback chain is now strict: a generator that claims a project and then fails aborts the run instead of silently falling through to a lower-quality scanner.

Run sbomify-action in an air-gapped CI environment without fetching external runtime bundles.

```yaml
- uses: sbomify/sbomify-action@v26.8.0
  env:
    SBOMIFY_TOKEN: ${{ secrets.SBOMIFY_TOKEN }}
    SBOMIFY_FETCH_RUNTIMES: '0'
```

— Names the env var, repo, and exact size drop; fallback mechanism explained

launch-20260825-ae80c5f8

---

04 **CRA patch SLA targets and KPI dashboard** NEW

85

New CRA patch SLA targets, defaulting to Critical 7 / High 30 / Medium 90 days, gate five computable KPIs: mean time to triage, mean time to remediate, open criticals with KEV split, patch SLA compliance, and VEX publication latency.

Set patch SLA targets by severity so CRA KPIs are computable from day one without opening workspace settings manually.

> 📍 In the workspace Settings, go to the CRA section and set patch SLA targets: Critical: 7 days, High: 30 days, Medium: 90 days. Once saved, the five KPIs (mean

time to triage, mean time to remediate, open criticals with KEV split, patch SLA compliance, VEX publication latency) become visible on the findings dashboard.

— Exact defaults and five named KPIs, with a UI path to set them `launch-20260825-b3c23b90`

05 **`risk_accepted` finding state** NEW 80

New `risk_accepted` finding state requires a business justification, named risk owner, and expiry date; it publishes downstream as `exploitable` with a `will_not_fix` response in CycloneDX rather than `not_affected`, to avoid false claims.

— Precise required fields and downstream CycloneDX mapping named `launch-20260825-b3c23b90`

06 **CSV exports for inventory, licenses, findings, vulnerabilities**

NEW 80

Four new CSV exports covering inventory, licenses, findings, and vulnerabilities, scoped to workspace, product, release, or component, reading the same latest-per-component data and live VEX suppressions as the dashboard.

Export all findings for a specific product release to CSV for auditor or procurement review, matching exactly what the dashboard shows.

📍 In the product release view, open the Exports section and download the Findings CSV. The export applies live VEX suppressions and latest-per-component data, so the row counts match the dashboard beside it. Unreadable artifacts appear as explicit rows rather than being silently omitted.

— Four formats and scopes named with a concrete download path `launch-20260825-b3c23b90`

07 **New ecosystem lock-file and manifest support** NEW 75

Adds first-class lock-file and manifest support for Gradle (`gradle.lockfile`), Swift (`Package.swift`), PHP (`composer.json` via a proper Composer bundle), Rust crates with only `Cargo.toml` , .NET project files, Haskell, Erlang, and Clojure. Manifests listing version ranges now read the adjacent lock file to record the versions actually shipped.

— Every named file format listed, but no CLI flag to trigger `launch-20260825-ae80c5f8`

08 **BSI-compliant security.txt generation** NEW 75

`security.txt` generation follows BSI TR-03183-3 shape with separate CSIRT contact and report-page slots, an optional CSAF pointer, and an `Expires` field clamped to at most one year; it shares data with the CRA wizard.

— Names the standard, exact field, and clamp value `launch-20260825-b3c23b90`

09 **CRA wizard EU Authorized Representative capture** NEW

70

CRA wizard now captures EU Authorized Representative details and mandate per Article 22/25, rendering them on the Declaration of Conformity; missing mandatory fields are surfaced loudly rather than silently omitted.

— Cites exact regulation articles and output document `launch-20260825-b3c23b90`

10 **Trust Center redesign with advisory table and faceted filters** IMPROVED

70

Trust Center redesigned with a filter sidebar, dense advisory table with per-product version rows rendering version ranges as comparison strings (e.g. `>= 11.2.0, < 11.2.4-h20`), live faceting, and a JavaScript-free plain form fallback.

— Concrete UI elements and an example version-range string named `launch-20260825-b3c23b90`

11 **End-of-life readiness check** NEW

65

End-of-life readiness check surfaces unpatched critical and high findings, drafts an EOL announcement through the advisory channel, and pins the final SBOM and VEX to a versioned release instead of a floating latest pointer.

— Three concrete behaviours described, no exact navigation given `launch-20260825-b3c23b90`

12 **Malicious packages reclassified as distinct category** IMPROVED

60

Malicious packages are now a distinct structural category, reclassified at read time across all historical scans without rescanning, and no longer fall through to a default medium severity.

— Clear before/after behaviour, no UI or API path given `launch-20260825-b3c23b90`

THINNER COVERAGE BELOW

13 **Install-free run via pipx or uvx** NEW

55

Supports running directly from PyPI with no install step via `pipx run sbomify-action` or `uvx sbomify-action`.

— Two exact runnable commands given, little else explained `launch-20260825-ae80c5f8`

14 **Explicit states for empty or unreadable scans** IMPROVED

55

Empty or unreadable scan results are now three distinct states — recognized nothing, scanner error, unreadable spec version — none of which render as a clean bill of health.

— Three named states, no UI location given for viewing them `launch-20260825-b3c23b90`

15 **Export retention floor per Article 13(15)** NEW

55

Export packages now pin a retention floor at creation per Article 13(15), so a later product change cannot silently shorten a previously exported compliance record.

— Names regulation clause and guarantee, no UI path given launch-20260825-b3c23b90

16 **Cross-run enrichment caching** IMPROVED 53

Enrichment lookups across all twelve sources are now cached across runs; transient failures and 429s are not persisted, and negative results expire in one day rather than one month.

— Names cache scope and expiry change, no config surface given launch-20260825-ae80c5f8

17 **Expanded vulnerability data: EPSS, CVSS 4.0, ENISA EUVD** IMPROVED 50

EPSS scores and CVSS 4.0 are now captured from both scanner paths; ENISA EUVD exploited flag appears alongside CISA KEV.

— Names four data sources but no way to act on them directly launch-20260825-b3c23b90

18 **Git submodule lock-file discovery** NEW 45

Discovery now detects lock files inside git submodules (via `.gitmodules` or vendored clones), annotates them clearly, and leaves them deselected by default with an explanation.

— Describes behaviour but no flag or command to invoke it launch-20260825-ae80c5f8

19 **Navigation spotlight search** IMPROVED 45

Navbar search upgraded to a navigation spotlight allowing jump-to for any product, component, or page.

— Bare description of a UI change, no shortcut or path given launch-20260825-b3c23b90

20 **Release-tag versioning for action builds** IMPROVED 35

Component version is now sourced from the release tag when the workflow runs on a tag; pre-releases are marked as pre-releases when published to sbomify.

— Thin description of a build-time behaviour change launch-20260825-ae80c5f8

└──▷ ALSO FROM THESE RELEASES

Bootstrap a workflow for your repository interactively without writing YAML by hand.

```
$ docker run --rm -it \
  -v "$(pwd):/github/workspace" \
  -w /github/workspace \
  -e SBOMIFY_TOKEN=$SBOMIFY_TOKEN \
  ghcr.io/sbomify/sbomify-action \
  sbomify-action wizard
```

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ 3 RELEASES · 2026-08-26 → 2026-08-27 NOTES ↗

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical expanded its Privileged Access Management module with PostgreSQL account discovery, IAM authentication, and dedicated secret folders, while also shipping certificate profile creation inside applications, root encryption key rotation, KMS key delete protection, RFC 8693 token exchange support, and change/approval policies for machine identities.

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

01 **Delete protection for KMS keys** NEW 70

Adds delete protection for KMS keys, preventing accidental key removal. Enable it in the Infisical console under project > Key Management, select a key, open its settings, and enable Delete Protection.

Protect a KMS key from accidental deletion in environments where key loss would be catastrophic.

> 📍 In the Infisical console, go to your project › Key Management, select a key, open its settings, and enable Delete Protection.

— Includes exact UI steps to enable the feature v0.163.0

02 **Helm chart now requires siteUrl config** BREAKING 60

The `infisical-nkp` Helm chart now requires `config.siteUrl` to be set in its configuration.

— Names exact chart and config key requiring migration v0.164.1

THINNER COVERAGE BELOW

03 **PostgreSQL support and secret organization in PAM** NEW 55

Adds PostgreSQL account discovery for Privileged Access Management (PAM), adds IAM authentication for Postgres in PAM enabling passwordless database credential management, and adds dedicated folder views for organizing PAM secrets.

— Names PAM, Postgres, IAM auth and folder views but no exact UI path v0.163.0 v0.164.1

04 **Dynamic secrets provider forms migrated to v3 UI** IMPROVED

45

Migrates dynamic secrets provider forms to an updated v3 UI, covering identity and access provider forms, data-service and protocol provider forms, and relational, warehouse, and managed-store provider forms.

— Names every provider category migrated but no exact navigation v0.163.0 v0.164.1

05 **RFC 8693 token exchange grant support** NEW

45

Adds RFC 8693 token exchange grant support for OAuth/OIDC flows.

— Names the exact standard but no usage steps given v0.163.0

06 **Certificate manager UI renewal enhancements** IMPROVED

45

Allows editing certificate details during manual certificate renewals, and shows renewed certificates with a Renewed status indicator in the certificate manager UI.

— Names both UI changes but no exact navigation path v0.163.0

07 **Gateway pool load balancing and clock skew handling** IMPROVED

40

Adds load balancing support for gateway pools and adds clock skew handling for gateway and relay connections.

— Describes both changes but no mechanism detail v0.163.0 v0.164.1

08 **Certificate profile creation inside applications** NEW

40

Allows creation of certificate profiles from directly inside applications, streamlining PKI workflows without leaving the app context.

— Explains the workflow benefit but no exact steps or endpoint v0.164.0

09 **Change/approval policies for machine identities** IMPROVED

40

Extends change/approval policies to apply to machine identities, not just human users.

— States scope expansion but no configuration detail v0.163.0

10 **Root encryption key rotation** NEW

35

Implements root encryption key rotation, enabling operators to rotate the platform's root encryption key.

— States the capability but not the mechanism or command v0.164.0

11 **Cross-project secret sharing gated by settings** IMPROVED

35

Gates cross-project secret sharing actions by settings, giving admins finer control over which projects can share secrets.

– Names the control but not the specific setting or location v0.164.0

12 **Project tag editing in project settings** NEW

30

Adds project tag editing under project settings.

– Bare feature name with minimal elaboration v0.163.0

13 **Project roles access control UI redesign** IMPROVED

25

Migrates the project roles access control UI to the v3 design.

– Bare statement of a UI migration with no further detail v0.163.0

WAS THIS USEFUL?  

◆ Web Application Security

slowhttptest

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

Application Layer DoS attack simulator

> slowhttptest v1.10.0 adds support for testing authenticated targets via mTLS client certificates and custom request headers, plus a build flag for non-standard OpenSSL installs.

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

01 **--with-openssl=DIR configure flag** NEW 65

Adds `--with-openssl=DIR` configure flag for builds where OpenSSL is installed in a non-standard path.

— Named build flag gives a concrete, runnable configure option v1.10.0

02 **mTLS and custom headers for authenticated targets** NEW 60

Adds mTLS support, enabling client certificates when testing targets that require mutual TLS authentication, and custom request header injection, allowing arbitrary headers to be sent with slow HTTP test payloads.

— Describes two new capabilities but no flag names or config given v1.10.0

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ 1 RELEASE · 2026-08-27 BLOG ↗

DataDome is a bot management and API protection platform that defends against automated attacks, web scraping, and credential abuse.

> DataDome expanded its API reference with new detection request fields and launched a hosted MCP server that lets AI agents query Trend Reports data in natural language.

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

01 New detection request fields in API reference NEW

60

The DataDome API reference gains several new fields for detection requests: `AcceptLanguage` for passing the request's Accept-Language header to the detection engine, `Via` for forwarding proxy chain information, `SecCHUAFullVersionList` for passing the `Sec-CH-UA-Full-Version-List` client hint header, and `CustomFieldString1`, `CustomFieldString2`, and `CustomFieldString3` for sending arbitrary custom string data alongside detection requests.

— Names all six fields but no usage context or limits given. api-docs-c11305f974b4

THINNER COVERAGE BELOW

02 Hosted MCP server for Trend Reports NEW

58

A new hosted MCP server exposes DataDome Trend Reports data to MCP-compatible AI agents and agentic workflows, authenticated via existing DataDome credentials with no custom integration or infrastructure required. It lets agents query automated traffic trends, threat types, scraping activity, and credential-stuffing patterns in natural language from tools like Slack-connected assistants, enterprise copilots, or custom agent pipelines.

Connect an MCP-compatible AI agent to DataDome's hosted MCP endpoint to query Trend Reports in natural language from within existing workflows.

> 📍 In your MCP-compatible AI agent or tool, configure the server connection to DataDome's hosted MCP endpoint and authenticate using your existing DataDome credentials.

— Explains mechanism and use cases but endpoint URL not given. launch-20260827-e988d370

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

ALEAPP

Feedback ⓘ 1 RELEASE · 2026-08-27 NOTES ↗ CODE ↗

Android Logs Events And Protobuf Parser

> ALEAPP v2026.3.2 adds parsers for 35+ new Android apps and significantly deepens support for TeleGuard, Threema, Uber, Ryanair and Truecaller, alongside Gmail parsing and reporting path fixes.

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

01 **Threema database parsing via sqlcipher3** NEW 75

Adds Threema account, contacts, and messages parsing, requiring a new `sqlcipher3` version 0.6.2 dependency to decrypt Threema's encrypted databases.

— Names exact dependency and version needed to run it v2026.3.2

02 **Truecaller call and contact parsing** NEW 75

Adds Truecaller call history, contacts, SMS senders, call cache, and settings parsed from `tc.db`, `insights.db`, and `calling-cache.db`.

— Names exact source database files parsed v2026.3.2

03 **Ryanair boarding pass decryption** NEW 70

Adds Ryanair boarding pass parsing with decryption using the key stored in the app's own preferences, plus bookings, flights, passengers, and sessions.

— Explains decryption key source, a concrete mechanism v2026.3.2

04 **Uber location and account artifacts** NEW 65

Adds Uber location artifacts with KML output for coordinates and timestamps, plus sessions, payment profiles, telemetry messages, cached images, and app settings.

— Names output format and artifact scope, no field-level detail v2026.3.2

05 **Gmail multi-account parsing improved** IMPROVED 65

Gmail parsing now reads every account store across all Android users instead of only the first, and Gmail bodies are now shown as readable text with a dedicated Links column.

— Describes before/after behaviour and new UI column v2026.3.2

06 **35+ new Android app parsers added** NEW 60

Adds parsers for over 35 new Android apps: Uber, Facebook, Threads, Twitch, Zoom, PayPal, Revolut, Outlook, Yahoo Mail, Adobe Acrobat Reader, Disney+, Amazon Prime Video, Pinterest, Hinge, NordVPN, Roblox, eBay, Etsy, Vinted, Booking.com, Ryanair,

DHL, CapCut, Shazam, Truecaller, SimpleX Chat, TeleGuard, Threema, Opera Browser, Samsung Bixby Routines, Spotify (6 artifacts), and Private Photo Vault.

— Broad list of app names but no per-app mechanism detail v2026.3.2

THINNER COVERAGE BELOW

07 **Reports use real device source paths** IMPROVED 50

Reports now use real device/source paths instead of examiner folder paths, with extraction-relative paths enforced library-wide.

— States the change but no config or migration steps given v2026.3.2

08 **TeleGuard artifact parsing added** NEW 40

Adds TeleGuard artifact coverage for calls, account, drafts, and downloads on Android.

— Names artifact types but no extraction mechanism v2026.3.2

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ 1 RELEASE · 2026-08-27 NOTES ↗ CODE ↗

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.3.2 adds parsers for 13 new apps, tightens cross-app data isolation, corrects reported source paths across 83+ modules, and ships a new CI guard against duplicate artifact path patterns.

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

01 **CI guard against duplicate case-only artifact paths** NEW 90

A new `admin/scripts/check_artifact_paths.py` CI guard, with `--self-test` and `--list` flags, fails when an artifact's `paths` tuple contains two patterns differing only in case; it is wired into the lint workflow.

– Exact script path and flags make this immediately runnable v2026.3.2

02 **Real device/source paths in reports** IMPROVED 70

iLEAPP now reports real device/source paths instead of the examiner's local folder structure, a change enforced across 83+ modules including Biome, CNA, and Privacy Accounting artifacts.

– Gives module count and names, a clear before/after fix v2026.3.2

THINNER COVERAGE BELOW

03 **13 new app parsers added** NEW 55

iLEAPP adds parsers for Google Calendar, Google Sheets, Google Docs, and Google Maps; Netflix (reporting only rows an examiner can act on, including cached stream segments with subtitle cues); Pinterest; Coinbase; Firefox; Grindr (including conversation link counts); Fitbit; KeepSafe (vault items, albums, and account/PIN security); and TeleGuard (calls, account, and drafts artifacts).

– Names every new app but no flags or commands to invoke them v2026.3.2

04 **Container-scoping enforcement for cross-app artifacts** IMPROVED 55

iLEAPP adds container-scoping enforcement for Pinterest, Tinder, and Google Drive/Sheets artifacts, preventing cross-app data bleed between them.

– Names affected apps but mechanism of enforcement not detailed v2026.3.2

WAS THIS USEFUL? 👍 👎

ElcomSoft

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

Try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

> Elcomsoft System Recovery 8.38 streamlines BitLocker TPM exploitation, adds broad browser artefact extraction, and automates checksum logging for disk images.

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

01 **Bootable USB for BitLocker TPM exploit** NEW 70

The installer can now create a separate bootable USB drive preloaded with a BitLocker TPM exploit, supporting `YellowKey` and `GreatXML`, turning what was previously a manual multi-step process into a guided setup click.

— Names exploit codenames and the workflow change, but no exact command shown

launch-20260826-274190ce

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

CHIPSEC

Feedback ⓘ 1 RELEASE · 2026-08-27 NOTES ↗ CODE ↗

Platform Security Assessment Framework

> CHIPSEC 2.0.8 adds HOB parsing, MSR scope validation, and enabled-state filtering for controls, alongside several installation and packaging fixes.

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

01 **Installation and packaging fixes** IMPROVED 65

Configuration files nested below `cfg/<vendor>/` are now correctly installed during package setup; the Linux kernel driver is now loaded via `modprobe` before searching for the driver file, improving portability of installed builds; and the manual is now installed to a path relative to the install prefix rather than a hardcoded location.

– Names exact paths and mechanism for each packaging fix 2.0.8

02 **is_enabled field and filter_enabled for controls** NEW 60

Adds an `is_enabled` field to controls, enabling `filter_enabled` filtering in object lists of controls.

– Names exact field and filter but no usage example 2.0.8

THINNER COVERAGE BELOW

03 **MSR scope correctness checks** NEW 40

Adds MSR scope correctness checks to catch misconfigured MSR definitions.

– States purpose but not how checks are run 2.0.8

04 **Exclusive fixed addresses in configuration helper** IMPROVED 40

The configuration helper now supports exclusive fixed addresses for platform register definitions.

– Describes change but not config syntax or example 2.0.8

05 **HOB parsing and inspection** NEW 30

Adds HOB (Hand-Off Block) parsing and inspection features to CHIPSEC.

– Named capability but no mechanism or usage detail given 2.0.8

06 **Enhanced logger messaging** IMPROVED 15

Enhanced logger messaging provides additional detail to aid debugging efforts.

— Generic description with no specifics on what changed `2.0.8`

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Abnormal AI

Feedback ⓘ 4 RELEASES · 2026-08-26 BLOG ↗

Abnormal AI is a behavioral AI security platform that protects email, identity, and AI—and stops insider threats. Trusted by 4,500+ organizations.

> Abnormal AI shipped three major additions this window: a no-cost Control Center for inbound email adding Custom AI Models and Custom Rules, an Early Access Email DLP Rules engine with an AI Triage Agent for outbound policy enforcement, and an upgraded AI Phishing Coach with Precision Training, AI-generated content creators, QR code simulations, and in-thread follow-up.

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

01 **Control Center for inbound email security** NEW 80

Control Center is a no-cost capability of Inbound Email Security offering two detection layers. Custom AI Models (Generally Available) let analysts describe a detection pattern in natural language to create a dedicated model, which is backtested against recent organizational mail and can run in Passive mode before full enforcement. Custom Rules (Early Access) let teams build organization-specific enforcement using more than 50 message attributes combined with `AND`, `OR`, and Except When logic — no proprietary query language required — with per-user or org-wide scoping, and can also be validated in Passive mode before going Active. Each Custom Rule and Custom AI Model assigns a verdict of Attack, Spam, or Safe, with Custom Rules taking precedence over Custom AI Models, which in turn override Abnormal's core AI detection; decision transparency shows which of the three layers made each detection decision.

When you want to block or allow email based on an organization-specific pattern without writing detection code, create a Custom Rule in Control Center and validate it against real examples before enabling it.

📍 In Abnormal, go to Control Center › Custom Rules › New Rule, specify conditions using available message attributes (sender domain, subject keywords, header values, etc.), test the rule against real message examples, and publish to enforce the policy.

— Names mechanism, modes, verdicts and precedence but only UI navigation given

launch-20260826-de891225 launch-20260826-af6414fb

02 **Email DLP Rules with AI Triage Agent** NEW 80

Email DLP Rules (Early Access) let security and compliance teams define, validate, and enforce outbound email policies via a rule builder supporting `regex`, phrase matching, metadata conditions, and Boolean logic — including AND/OR nesting, condition groups,

and exclusions — with prebuilt templates for data types like financial information and PII. Rules scan outbound content and supported attachments and can be scoped by department or workflow. The AI Triage Agent evaluates each flagged message against the rule's stated intent, sender-recipient relationship, and surrounding context to distinguish genuine violations from routine business traffic, auto-releasing benign messages and quarantining likely violations; it also generates a plain-language description of each rule's intent and editable exclusions at creation time, and rules can be validated against sample emails before live deployment. Every Triage Agent verdict is recorded in a new Outbound Log for audit visibility.

Review Triage Agent verdicts and reasoning for outbound policy matches to prioritize which flagged messages need analyst attention.

> 📍 In the Abnormal console, go to Outbound Log, filter by the relevant Email DLP Rule, and inspect the plain-language Triage Agent explanation alongside each flagged message verdict.

When you need to suppress a known-safe outbound pattern (e.g., invoices from your billing system) without manual SOC review, tune an Email DLP Rule exclusion so the AI Triage Agent auto-releases those messages.

> 📍 In Abnormal, go to Email DLP Rules › New Rule, define your policy using regex or phrase matching, add exclusions in plain language (e.g., 'exclude messages from billing@yourcompany.com'), and enable AI Triage Agent review to auto-release matches with benign business context.

— Detailed rule syntax and triage logic, but only UI steps shown `launch-20260826-af6414fb` `launch-20260826-7a61fbd3`

---

03 AI Phishing Coach: Precision Training, creators, QR sims, follow-up

NEW 80

AI Phishing Coach gained: Precision Training, which automatically adjusts simulation frequency and difficulty per employee based on behavior signals and response history, giving more frequent/harder simulations to elevated-risk employees and fewer to low-risk ones, plus reinforcement training when needed; Attack Creator, which generates a realistic simulation from a plain-language scenario description and stores it in the Attack Library; Campaign Creator, which schedules and sends ad hoc or baseline simulation campaigns from the Attack Library, including AI-generated attacks; Training Creator, which generates custom training video modules from a prompt and imagery, with module recommendations driven by Abnormal's observed threat patterns; QR Code Phishing Simulations, which train employees to treat QR codes with the same suspicion as suspicious URLs and render even in email clients that block images; and In-Thread Follow-Up, which re-engages employees who neither clicked nor reported a simulation by sending a follow-up inside the same email thread, mirroring real attacker tactics.

When a specific employee repeatedly ignores phishing simulations, use In-Thread Follow-Up in AI Phishing Coach to automatically re-engage them in context rather than waiting for the next scheduled campaign.

> 📍 In Abnormal, go to AI Phishing Coach › Campaign Creator, create or select a campaign with In-Thread Follow-Up enabled, then Precision Training will automatically increase simulation frequency for non-responsive employees and reduce it as behavior improves.

— Six named sub-features with mechanism, only UI path for use `launch-20260826-af6414fb` `launch-20260826-3d0f5330`

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Bubblewrap

Feedback ⓘ 1 RELEASE · 2026-08-26 NOTES ↗ CODE ↗

Low-level unprivileged sandboxing tool used by Flatpak and similar projects

> Bubblewrap v0.12.0 introduces a fail-open sandboxing mode and a build-time kernel assumption option, while dropping setuid support entirely in favor of unprivileged user namespaces.

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

01 `--not-a-security-boundary` flag for fail-open sandboxing

NEW 75

The new `--not-a-security-boundary` flag makes failures in sandbox setup steps (such as remounting a submount) non-fatal, enabling fail-open behavior for use cases that do not require a strict security boundary.

Run a sandbox for a use case where filesystem layout isolation matters but strict security enforcement is not required, so that non-fatal setup failures are tolerated rather than aborting.

```
$ bwrap --not-a-security-boundary --ro-bind /usr /usr --proc
/proc --dev /dev --unshare-pid bash
```

– Names exact flag and gives runnable example. v0.12.0

02 `assume_kernel` meson build option NEW

75

A new `assume_kernel` meson build option lets builders specify a minimum kernel version; specifying `5.6.0` or later disables the fallback implementation of openat2(RESOLVE_IN_ROOT) and removes backward compatibility code for older kernels, resulting in hard failures at runtime on unsupported kernels.

– Names build option and exact syscall/version, but no example. v0.12.0

THINNER COVERAGE BELOW

03 Setuid build mode removed BREAKING

52

Support for building a setuid bubblewrap has been removed; unprivileged user namespaces are now the sole supported mechanism. Deployments relying on the setuid build mode will no longer work and must migrate to unprivileged user namespaces.

– Clear before/after but no migration steps given. v0.12.0

└──▷ **BREAKING ON UPGRADE**

! Setuid bubblewrap support has been removed entirely; deployments relying on the setuid build mode will no longer work and must migrate to unprivileged user namespaces.

WAS THIS USEFUL?  

Alibaba Cloud OpenSandbox

Feedback ⓘ

3 RELEASES · 2026-08-26 NOTES ↗

Secure, Fast, and Extensible Sandbox runtime for AI agents.

> OpenSandbox's server v0.2.3 adds QEMU-based pause/resume for runc sandboxes, machine-identifiable pool-exhaustion signaling, and OSEP-0020 lifecycle hooks that now span both the server and execd 1.1.0, which itself ships PID-1 init mode, a pre-exec hardening pipeline (seccomp/Landlock), and a new eBPF audit image variant.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 **Sandbox lifecycle hooks (preStart/periodic)** NEW 90

Adds `CreateSandboxRequest.lifecycle` (`preStart` / `periodic` hook types per OSEP-0020), transported to Kubernetes workloads via the reserved `OPEN_SANDBOX_LIFECYCLE` env var; the Docker provider explicitly rejects lifecycle requests. Config persists at `$HOME/.execd/lifecycle.toml` with `EXECD_LIFECYCLE_CONFIG` as an exact override path; execd starts its HTTP server before `preStart` and only releases the user process after `preStart` succeeds, with bootstrap sync and preStart shutdown fail-closed under TERM/KILL watchdogs. The `preStart.timeoutSeconds` server and OpenAPI limit was raised from 300s to 10800s (3 hours); the `periodic` limit stays at 300s.

Configure a preStart hook and a periodic health check for a sandbox, pointing execd at a custom lifecycle config file.

```
$ EXECD_LIFECYCLE_CONFIG=/etc/sandbox/lifecycle.toml execd
```

— Names config path, env var, timeouts, and exact success/failure behavior. server/v0.2.3 docker/execd/1.1.0 docker/execd/v1.1.0

02 **Pool exhaustion signaling and allocation visibility** IMPROVED 90

Adds `kubernetes.pool_acquisition_timeout_seconds` config key (default 30s) so the server waits a bounded window on pool exhaustion, then returns HTTP 429 with `KUBERNETES::POOL_CAPACITY_EXHAUSTED` and `Retry-After: 5` instead of a generic 504. Pool exhaustion is now machine-identifiable via a controller-published `PoolAllocationPending` / `PoolCapacityExhausted` condition, and `getSandbox` surfaces an optional `Sandbox.allocation` block (`mode: pool` , `poolRef` , `state: allocated`) for confirmed Kubernetes pool allocations.

Tune how long the server waits for pool allocation before returning a machine-identifiable 429, and increase the egress sidecar readiness window for slower Docker hosts.

```toml
[kubernetes]
pool_acquisition_timeout_seconds = 60

[egress]
timeout_seconds = 60
```

— Exact status codes, headers, condition names, and config key given. `server/v0.2.3`

03 **execd init-mode as PID 1** NEW 85

Enables execd to run as PID 1 via `EXECD_INIT=1` (or server-side `runtime.execd_run_as_init`), providing a single reaper with signal forwarding (TERM/HUP/USR1/USR2/WINCH), entrypoint-owned exit-code propagation, and a subreaper fallback for the Pool path; classic topology is unchanged and this is off by default.

Run execd as PID 1 in a Docker container to get proper signal forwarding and exit-code propagation for sandboxed workloads.

```
$ docker run -e EXECD_INIT=1 opensandbox/execd:v1.1.0
```

Enable init-mode to get proper PID-1 signal forwarding and exit-code propagation in a Docker sandbox container.

```
$ docker run -e EXECD_INIT=1 opensandbox/execd:v1.1.0
```

— Runnable docker command plus exact signal and fallback behavior given.

`docker/execd/1.1.0` `docker/execd/v1.1.0`

04 **Pre-exec hardening pipeline in execd** NEW 78

Adds a new `opensandbox-launcher` pre-exec hardening pipeline (env strip → KEEPCAPS → bounding-set trim → `no_new_privs` → identity drop → seccomp → execve), with every layer degrading fail-open. A new `[landlock]` layer applies a filesystem allowlist, failing open below ABI/kernel 5.13 and reporting `unsupported` state. Per-layer state is reported in a new `hardening` object on `GET /v1/isolated/capabilities`.

Inspect the per-layer hardening state of a running sandbox to verify seccomp, no_new_privs, and Landlock are active.

```
$ curl -s GET https://<host>/v1/isolated/capabilities | jq '.hardening'
```

Inspect per-layer hardening state (seccomp, landlock, no_new_privs, etc.) reported by the new `hardening` object to verify the pre-exec floor is active.

```
$ curl -s https://<sandbox-host>/v1/isolated/capabilities | jq '.hardening'
```

— Named endpoint and full pipeline order, but check is read-only. `docker/execd/1.1.0` `docker/execd/v1.1.0`

---

05 **Secure-access key injection from Kubernetes Secrets** NEW

78

Accepts `OPENSANDBOX_SECURE_ACCESS_KEYS` and `OPENSANDBOX_SECURE_ACCESS_ACTIVE_KEY` env overrides for secure-access signing keys; the Helm chart adds `server.gateway.secureAccess.existingSecret` (mutually exclusive with plaintext `keys`) so GitOps flows never commit key material.

Keep signing-key material out of Git by pointing the Helm chart at an existing Kubernetes Secret instead of committing plaintext keys.

```yaml
server:
  gateway:
    secureAccess:
      existingSecret: opensandbox-signing-keys
```

— Named env vars and Helm key with a working values.yaml example. `server/v0.2.3`

---

06 **QEMU VMState pause/resume for runc sandboxes** NEW 78

Adds opt-in QEMU VMState pause/resume for runc-based sandboxes: capture of guest memory, vCPU, and device state via QMP migration, stored as separate OCI images keyed by immutable SHA-256 digests; supports pooled BatchSandbox pause, detachment, and standalone restoration.

— Mechanism and storage model detailed, but no invocation command given. `server/v0.2.3`

---

07 **Fleet-wide Docker sandbox env and bind mounts** NEW

75

New `[docker]` config fields `sandbox_env` and `sandbox_binds` apply fleet-wide environment variables and host bind mounts to every Docker sandbox container; request `env` wins over `sandbox_env`, and config binds precede request-derived binds.

Apply fleet-wide environment variables and a host bind mount to every Docker sandbox so operators can inject shared config without touching each request.

```toml
[docker]
sandbox_env = { "LOG_LEVEL" = "info", "REGION" = "us-east-1" }
sandbox_binds = ["/etc/ssl/certs:/etc/ssl/certs:ro"]
```

— Named config keys with precedence rules and a runnable example. `server/v0.2.3`

---

08 **Scoped Diagnostics API responses** IMPROVED 68

Diagnostics API now honors scopes: `logs: container|all` and `events: runtime|lifecycle|all`, returning 400 `DIAGNOSTICS_SCOPE_UNSUPPORTED` for unknown scopes; deprecated no-`scope` responses still work but now include a `Deprecation` header.

— Named scope values and error code but no endpoint path given. `server/v0.2.3`

---

09 **OTLP HTTP request metrics** NEW 65

Adds generic HTTP request metrics via OTLP: a thin ASGI middleware exports `server.http.request.duration` (ms) with `http_method`, `http_route`, and `http_status_code` attributes; no-op when `[otel]` is disabled.

— Named metric and attributes but no example query or dashboard. `server/v0.2.3`

---

10 **Egress timeout and extra-port configuration** NEW 62

New `[egress] timeout_seconds` config key (default 30s; non-positive values rejected) lets Docker deployments tune how long the server waits for the sidecar health endpoint. An experimental `OPENSANDBOX_EGRESS_MITMPROXY_EXTRA_PORTS` env var appends extra TCP ports to the always-on 80/443 interception list for the egress transparent MITM, capped at 15 ports.

Tune how long the server waits for pool allocation before returning a machine-identifiable 429, and increase the egress sidecar readiness window for slower Docker hosts.

```toml
[kubernetes]
pool_acquisition_timeout_seconds = 60

[egress]
timeout_seconds = 60
```

— Two named config surfaces with defaults and a cap value. `server/v0.2.3`

11 **eBPF audit variant for execd** NEW 60

Introduces an `execd-ebpf` image variant (CGO + cilium/ebpf) that hooks exec/connect/privilege events via CO-RE, scoped to the sandbox cgroup, and emits a rotating JSONL audit log; requires kernel ≥5.10 with BTF.

— Named image variant and kernel requirement but no usage example. `docker/execd/1.1.0`

THINNER COVERAGE BELOW

12 **Egress sandbox ID no longer client-settable** BREAKING 55

`OPENSANDBOX_EGRESS_SANDBOX_ID` is now server-injected and removed from the env allowlist; requests that pass it in `env` receive HTTP 400 instead of a user-spoofable value.

— Names exact env var and new error response for callers. `server/v0.2.3`

13 **Startup tenant namespace validation** NEW 50

Tenant namespaces are now validated at startup (OSEP-0014): the server fail-fast verifies every tenant namespace exists and is accessible, aggregating all failures into a single fatal error before serving traffic.

— Behavior described but no config key or command shown. `server/v0.2.3`

14 **OTLP endpoint env vars for egress auto-allow** IMPROVED 45

OTLP egress auto-allow now honors `OTEL_EXPORTER_OTLP_METRICS_ENDPOINT` and `OTEL_EXPORTER_OTLP_ENDPOINT` environment variables in shared components.

— Names two env vars but no mechanism or example. `docker/execd/1.1.0`

15 **FastPath v2 gRPC client (groundwork)** NEW 40

Adds a FastPath v2 gRPC client wrapping all FastPathService RPCs with typed error mapping (`FastPathNotFound` → 404, etc.) as groundwork for a fast-sandbox fleets backend (OSEP-0007); not yet wired as an active runtime.

— Explicitly groundwork with no active usage path yet. server/v0.2.3

└──▷ BREAKING ON UPGRADE

! `OPENSANDBOX_EGRESS_SANDBOX_ID` is now server-injected and removed from the env allowlist; requests that pass it in `env` receive HTTP 400 instead of a user-spoofable value.

! The Kubernetes provider transports lifecycle hooks via the internal `OPEN_SANDBOX_LIFECYCLE` environment variable, which is stripped from user code; the Docker provider rejects lifecycle requests entirely — workloads relying on Docker-side lifecycle hooks will receive errors.

WAS THIS USEFUL? 👍 👎

Tencent Cloud CubeSandbox

Feedback ⓘ

1 RELEASE · 2026-08-28 NOTES ↗

Instant, Concurrent, Secure & Lightweight Sandbox for AI Agents.

> CubeSandbox v0.7.0 ships cross-node pause/resume with S3-backed snapshot cloning, splits node management into a new multi-replica CubeOps control plane (breaking for cubemastercli scripts), refactors the network subsystem for lower sandbox-creation latency, and brings the Go and Node SDKs to feature parity with Python.

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

01 CubeOps control plane replacing CubeMaster node management

75

BREAKING

Introduces the `CubeOps` module with multi-replica deployment support, taking over node management from CubeMaster, and ships the `cubeopscli` tool for node management, replacing the `cubemastercli` node subcommands. Existing scripts using `cubemastercli` node subcommands will break.

— Names the new tool and the breaking change precisely v0.7.0

02 Cross-node pause/resume with S3-backed snapshot cloning NEW

65

Introduces cross-node pause/resume and snapshot cloning with an S3 backend (preview), backed by the new `CubeS3lvol` COW engine for remote snapshot storage on S3 object storage. MinIO is deployed by default as the S3 backend for this feature and for the volume feature, with the option to specify a custom S3 backend.

— Names the COW engine and backend but no config keys given v0.7.0

03 Node SDK Sandbox.setTimeout with NEVER_TIMEOUT NEW

65

Adds Sandbox.setTimeout(timeout) to the Node SDK with support for `NEVER_TIMEOUT`.

Set a Node SDK sandbox to never time out for long-running interactive sessions.

```javascript
await sandbox.setTimeout(Sandbox.NEVER_TIMEOUT);
```

— Named method plus runnable code example provided v0.7.0

04 **Python SDK Sandbox.create distribution_scope placement control**

NEW 65

Adds Sandbox.create(distribution_scope=...) to the Python SDK to explicitly specify the nodes/zones a sandbox can be placed on.

Pin a Python SDK sandbox to a specific zone or set of nodes to control placement in multi-region or compliance-sensitive deployments.

```python
sandbox = Sandbox.create(distribution_scope='zone-us-east-1')
```

— Named parameter plus runnable code example provided v0.7.0

THINNER COVERAGE BELOW

05 **Dynamic network policy updates, CubeEgress L7 forwarding, and CubeVS MAC learning**

NEW 55

Supports dynamically updating network policies for running sandboxes, adds support for `CubeEgress` L7 forwarding rules with custom ports, and adds CubeVS MAC address learning within the same subnet to avoid black-holed traffic in non-hairpin networks.

— Names three network components but no exact config syntax v0.7.0

06 **Separate kernel and guest image release artifacts** NEW

55

Publishes kernel and guest images as separate artifacts under dedicated `kernel-release-*` and `guest-image-*` releases.

— Names exact release artifact tags, but sparse context v0.7.0

07 **Component multi-versioning for templates and snapshots** NEW

50

Templates and snapshots now record their component dependency versions, allowing historical templates and snapshots to remain usable after compute-node component upgrades.

— Clear mechanism but no command or config to act on v0.7.0

08 **NetworkAgent merged into Cubelet** IMPROVED 50

NetworkAgent is now fully merged into Cubelet, reducing inter-process RPC calls during sandbox creation and preventing network rules from leaking across sandboxes.

— Mechanism and benefit stated, no user-facing action v0.7.0

09 **Node lifecycle management: deletion, isolation, and history**

NEW

50

Adds node deletion support so a node can be removed from the cluster and excluded from scheduling, plus node isolation/unisolation and node operation history viewing in the Web UI.

— Names Web UI surface but no exact navigation steps v0.7.0

10 **CubeProxy custom management port and gRPC protocol** NEW

50

Adds `CubeProxy` support for a custom management port and the gRPC protocol.

— Names the component and protocol, no config example v0.7.0

11 **Go SDK Files.ForUser for per-user isolation** NEW 50

Adds `Files.ForUser` to the Go SDK to isolate file access by user identity.

— Names exact API method, description otherwise thin v0.7.0

12 **Volume CRUD and volumeMounts parity in Go/Node SDKs** NEW

50

Adds Volume CRUD and `volumeMounts` to the Go and Node SDKs, bringing all three languages to parity with Python.

— Named field and scope, no usage syntax shown v0.7.0

13 **eBPF network policy provisioning latency optimization**

IMPROVED

45

Optimizes the eBPF network policy provisioning path to significantly reduce sandbox creation latency for sandboxes with network rules.

— States goal without numbers or a config surface v0.7.0

14 **Template alias support across Go, Node, Python SDKs** NEW

45

Adds full template alias support, including build-time aliases and alias management for existing templates, across the Go, Node, and Python SDKs.

— Describes scope across SDKs, no exact API names v0.7.0

15 **HTTP private registry image pulling for templates** NEW

35

Supports pulling images from HTTP (private) image registries to build templates.

— Thin one-line capability description, no config details v0.7.0

16 **Configurable Cubelet log rotation policy** NEW 30

Adds configurable log rotation policy for Cubelet via the configuration file.

— Bare mention of config file, no key names given v0.7.0

└──▷ **BREAKING ON UPGRADE**

! Node management subcommands previously in `cubemastercli` are replaced by `cubeopscli`; existing scripts using `cubemastercli` node subcommands will break.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗ CODE ↗

Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment.

> Prowler 5.40.0 introduces an Alerts feature with Slack and email delivery for findings, ships 15 new security checks across seven providers, and adds the NCSC Cyber Essentials 3.3 compliance framework alongside smaller UI improvements to Lighthouse AI and provider tracking.

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

01 **15 new checks across 7 providers** NEW

85

AWS: `ecr_repository_image_no_secrets` scans ECR image config and filesystem layers for hardcoded secrets; `bedrock_guardrail_contextual_grounding_filter_enabled` verifies Bedrock guardrails enable contextual grounding filters; `bedrock_custom_model_encrypted_with_cmk` and `bedrock_knowledge_base_encrypted_with_cmk` verify Bedrock models and knowledge-base storage use customer-managed KMS keys; `bedrock_agent_role_not_shared_across_agents` verifies each Bedrock agent has a dedicated execution role; `rolesanywhere_profile_restricts_session_permissions` flags IAM Roles Anywhere profiles referencing an administrative role without a scoping session or managed policy. GCP: `iam_workload_identity_pool_provider_attribute_condition` flags Workload Identity Federation providers trusting a multi-tenant issuer without an attribute condition. GitHub: `organization_default_workflow_permissions_read_only` and `repository_default_workflow_permissions_read_only` verify workflows get a read-only default `GITHUB_TOKEN`; `organization_actions_pull_request_approval_disabled` verifies orgs block Actions from creating and approving pull requests. Microsoft 365: `defender_domain_dmarc_records_published` verifies Exchange Online domains publish an enforcing DMARC policy (`p=quarantine` or `p=reject`). Alibaba Cloud: `oss_bucket_versioning_enabled` and `oss_bucket_server_side_encryption_enabled` check OSS bucket versioning and default AES256/KMS encryption. Huawei Cloud: `vpc_security_group_open_egress` flags security groups allowing open egress to the internet. STACKIT: `ske_cluster_no_public_endpoint` flags SKE clusters whose Kubernetes API endpoint is publicly reachable because the ACL extension is disabled or its allowed CIDR list contains `0.0.0.0/0` or `::/0`.

— Every check named verbatim with its exact behaviour 5.40.0

02 **Alerts for findings with Slack and email delivery** NEW

75

New Alerts feature lets you create alerts from Prowler Cloud findings and deliver them to email recipients and Slack channels, triggered after scans or on a daily digest schedule. Connect a Slack workspace from the Integrations page, authorize destination channels, and select them in the alert modal's Destination channels selector alongside the existing email Recipients selector; the alerts list shows both in a unified Destinations column. Available in Prowler Cloud and Prowler Private Cloud.

— Full mechanism and UI path given, no exact command 5.40.0

03 **NCSC Cyber Essentials 3.3 compliance framework** NEW

70

Adds NCSC Cyber Essentials 3.3 (Requirements for IT Infrastructure v3.3, April 2026) as a universal compliance framework with 28 requirements across five control themes: Firewalls, Secure Configuration, Security Update Management, User Access Control, and Malware Protection — with Azure provider coverage for 16 requirements.

— Names version, structure, and provider coverage but no scan command 5.40.0

THINNER COVERAGE BELOW

04 **Imported provider indicator in providers table** NEW

45

The Providers table now shows an Imported provider indicator for providers whose findings were uploaded via the Prowler CLI, distinguishing them from connected providers at a glance.

— Names the UI element and its purpose, no deeper mechanism 5.40.0

05 **Feedback on Lighthouse AI answers** NEW

35

Lighthouse AI answers now support per-answer thumbs-up/thumbs-down feedback with an optional free-text description field, collected directly in the chat.

— Describes the mechanism but stays high-level 5.40.0

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

cilium-cli

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

CLI to install, manage & troubleshoot Kubernetes clusters running Cilium

> cilium-cli v0.20.0 adds a reservedRanges field for multi-pool IPAM exclusions and HTTP CONNECT support in the Envoy integration.

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

01 **Reserved ranges in multi-pool IPAM** NEW 65

Adds a `reservedRanges` field to `CiliumPodIPPools`, letting operators exclude specific address ranges from pool allocation in multi-pool IPAM.

— Names exact config field but no usage example v0.20.0

THINNER COVERAGE BELOW

02 **HTTP CONNECT support in Envoy integration** NEW 40

Adds HTTP CONNECT support to the Envoy integration, enabling tunneled proxy traffic through Cilium-managed Envoy.

— States capability but no config or command detail v0.20.0

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ 1 RELEASE · 2026-08-28 CHANGELOG ↗

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

> Huntress shipped per-identity ITDR escalations, a standalone Executive Summary PDF export for incident reports, 7 new Microsoft 365 controls for Managed ISPM (nearing 80 total), and added visibility into full IPv4 addresses and WSL activity.

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

01 **Per-identity escalations for Unexpected Country and VPN** NEW

65

Managed ITDR now generates a separate escalation, notification, and (if a PSA integration is configured) a separate PSA ticket for each identity affected by Unexpected Country or Unexpected VPN activity, rather than bundling multiple affected identities into one alert.

Review and tune per-identity escalation notification recipients and PSA ticket categories before multiple concurrent ITDR escalations flood the wrong queues.

> 📍 In the Huntress portal, go to Account Settings › Huntress Platform and Alert Notifications, then review each category (Unexpected Country, Unexpected VPN) to confirm the correct recipients and PSA settings are configured for per-identity escalations.

— Names mechanism and trigger conditions; UI path given but no API/config key

changelog-20260828-32c0e05a

02 **7 new Microsoft 365 controls in Managed ISPM** NEW 65

Adds 7 new Microsoft 365 security controls to Managed ISPM, bringing the total to approaching 80: Anti-Malware, Anti-Spam, and Anti-Phishing policies checked against the Microsoft Standard baseline in Exchange Online Protection, plus Safe Links, Safe Attachments, Safe Documents, and spoof-protection Anti-Phishing policies in Microsoft Defender for Office 365.

— Names all 7 controls and baseline but gives no user action to take

changelog-20260828-32c0e05a

03 **Executive Summary PDF for ITDR incident reports** NEW

60

Every ITDR incident report now includes a plain-text Executive Summary at the top, with a standalone PDF export available via the `Export PDF` button in the portal.

Export a ready-to-share executive briefing from an active ITDR incident without rewriting technical findings for a client or insurer.

> 📍 In the Huntress portal, open the ITDR incident report, then click the `Export PDF` button at the top of the report to queue a standalone Executive Summary PDF download.

— Names the exact UI button and output format, no deeper mechanism

changelog-20260828-32c0e05a

THINNER COVERAGE BELOW

04 **Full IPv4 support via ipv4_addresses API field** NEW

30

Adds full IPv4 visibility through a new `ipv4_addresses` API field, per the release summary, though no further detail on scope or usage is given.

— Only the field name is given, no endpoint, schema, or usage changelog-20260828-32c0e05a

05 **WSL visibility** NEW

15

Adds visibility into Windows Subsystem for Linux (WSL) activity, mentioned only in the release summary with no further detail on what data is surfaced or where.

— Named only in passing, no mechanism, surface, or example given

changelog-20260828-32c0e05a

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ 1 RELEASE · 2026-08-27 BLOG ↗

Microsoft Defender is an antivirus and threat protection tool that scans for and removes malware, viruses, and security threats on Windows devices.

> Microsoft Defender expanded its managed detection service to cover third-party data sources and launched a new expert-led threat intelligence offering.

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

01 **MDR coverage for third-party data sources via Sentinel** NEW

60

Microsoft Defender Experts MDR now extends managed detection and response and threat hunting to third-party data sources ingested through Microsoft Sentinel — including Palo Alto Networks, AWS, and Okta — available under the MDR P2 tier.

— Names tier and integrated sources but no setup steps. 2026-08-27-8ddd7ba3

THINNER COVERAGE BELOW

02 **Defender Experts Threat Intelligence service** NEW 30

New Microsoft Defender Experts Threat Intelligence service delivers expert-led threat intelligence and actionable insights tailored to geography, industry, and risk profile.

— Describes purpose only, no mechanism or access details. 2026-08-27-8ddd7ba3

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Splunk Security Content

Feedback ⓘ 1 RELEASE · 2026-08-27

NOTES ↗

Splunk Security Content

> Splunk Security Content v6.5.0 adds 13 new detections spanning Python supply-chain attacks, Vidar stealer cloud credential theft, RoguePlanet/ShieldBreak Windows Defender abuse, and EDRSilencer evasion, alongside fidelity updates to 46 existing analytics.

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

01 Malicious Python Package Installation analytic story NEW

68

New analytic story with four detections covering Python supply-chain threats that execute code during install or persist across sessions: Python Network Traffic During Package Build, Python PTH File Creation During Package Installation, Python PYTHONPATH Modification During Package Installation, and Python Site Hooks Creation During Package Installation.

— Names all four detections and their persistence mechanism v6.5.0

02 RoguePlanet analytic story detections NEW

68

Six new detections added to the RoguePlanet analytic story targeting Windows Defender race-condition and ShieldBreak techniques: Windows Alternate Data Stream Created Over Local Share, `Windows Defender MpClient.dll Loaded by Non-Defender Process`, Windows Defender Threat Detected on Kernel Object Path, `Windows Error Report Created in ReportQueue Manually`, Windows Phantom DLL Created on Disk, and Windows Wermgr Spawning System Integrity Process.

— Lists all six detection names and the targeted technique v6.5.0

THINNER COVERAGE BELOW

03 EDRSilencer evasion detections via Windows Filtering Platform

NEW

59

Two new behavioral detections identifying attempts to block security processes from reaching management infrastructure: Windows EDRSilencer Custom Outbound Filter Added and Windows Filtering Platform Filter Added To Block EDR Process.

— Names both detections and the evasion mechanism v6.5.0

04 Vidar Stealer analytic story NEW

56

New analytic story anchored by Windows Cloud Sensitive File Read Access By Uncommon Process, detecting collection of Azure CLI profile metadata, saved credentials, and cookies to support account takeover and cloud reconnaissance.

— Names the detection and describes what it collects v6.5.0

05 **Fidelity updates to 46 existing analytics** IMPROVED

35

46 existing analytics updated across Windows, Linux, macOS, ESXi, and AWS Bedrock to improve fidelity across AI infrastructure abuse, credential access, browser data theft, process injection, persistence, privilege escalation, and destructive activity.

— No individual analytics or specific changes named v6.5.0

WAS THIS USEFUL? 👍 👎

Tenzir

Feedback ⓘ 2 RELEASES · 2026-08-18 → 2026-08-28 BLOG ↗

Tenzir is a data pipeline platform that ingests, transforms, and analyzes security telemetry and logs for threat detection.

> Tenzir's biggest addition is native RELP ingestion via a new `accept_relp` operator for rsyslog fleets, alongside native Splunk HEC ingestion, mergeable statistical models for drift detection, new TQL temporal operators, and a broad Explorer UI overhaul covering timestamp field selection, schema browsing, right-click filtering, and value-frequency previews.

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

01 **RELP ingestion via accept_relp operator** NEW 95

Adds `accept_relp` operator to receive syslog events over the Reliable Event Logging Protocol (RELP), preserving message boundaries (including embedded newlines), peer address, and RELP transaction ID alongside the payload. Delivery is at-least-once — Tenzir batches accepted RELP transactions and coalesces ordered acknowledgements, reaching ~0.5M events/second in internal raw transport benchmarks — and TLS/mutual TLS are supported for collectors crossing trust boundaries.

Configure rsyslog to forward logs to a Tenzir node over RELP with a disk-backed queue for durability across restarts.

```
$ global(workDirectory="/var/spool/rsyslog")
module(load="omrelp")
action(
  type="omrelp"
  target="tenzir.example.com"
  port="2514"
  action.resumeRetryCount="-1"
  queue.type="LinkedList"
  queue.filename="tenzir-relp"
  queue.saveOnShutdown="on"
)
```

Accept RELP-delivered syslog, parse each message, normalize via the OpenSSH package, and map to OCSF — then route failed logins to Splunk while archiving everything to Iceberg.

```
$ accept_relp "0.0.0.0:2514"
this = data.parse_syslog()
openssh::parse
openssh::ocsf::map
fork {
  where status_id == 2
  to_splunk "https://splunk.example.com:8088", hec_token =
secret("splunk-hec-token")
}
to_iceberg "security.ocsf", catalog =
"https://catalog.example.com", partition_by = [class_uid,
day(time)]
```

— Named operator, throughput number, and runnable rsyslog + pipeline config given

launch-20260818-9be4c3ae

02 **Explorer timestamp field selection for time ranges** NEW

80

Adds `_ts`, `ts`, `timestamp`, and `time` field selection to the Explorer time picker, plus custom field name entry, so time-range predicates target the correct timestamp field without overwriting unrelated pipeline filters.

Scope a time range to a non-default timestamp field when your events use a custom time column instead of the standard ingestion timestamp.

> 📍 In the Explorer time range picker, open Timestamp field selection and choose the field (e.g. your event's own time column) to use as the time axis.

— Named field options plus exact UI steps to scope a custom time field

Tenzir Platform v1.39.0: Explorer Usabi…

03 **Explorer stored data tab and schema browsing** NEW 75

Adds a Stored data tab in the Explorer that lets you browse node-wide schemas and their fields, event counts, and schema variants before running a pipeline, with `@name` filter injection and direct stored-event access.

Browse stored schemas before writing a pipeline to quickly identify which field carries timestamps in an unfamiliar dataset.

> 📍 In the Explorer, open the Stored data tab, search for a schema or field name, then click a schema to open its stored events directly.

Browse available node schemas in the Explorer to understand field structure before building a query.

> 📍 In the Explorer, open the schema panel for a connected node, then select a schema to inspect its fields and types.

— Names concrete tab and injection mechanism, with two navigation examples

Tenzir Platform v1.39.0: Explorer Usabi…

---

04 **Explorer Stream view display controls** NEW 70

Adds display controls to the Explorer Stream view: open the Fields panel to show or hide fields across schema variants, Shift-click a field to show it exclusively, and toggle value wrapping and field-name visibility from the Stream toolbar.

— Describes specific UI controls but no runnable example given

Tenzir Platform v1.39.0: Explorer Usabi…

---

05 **Right-click field context menus in Explorer** NEW 70

Adds right-click context menus on fields and values in Explorer Stream, Table, and Inspector views to build equality/comparison filters, string matches, timestamp windows, field selection, and field removal without writing TQL manually.

— Lists filter types produced but no step-by-step example

Tenzir Platform v1.39.0: Explorer Usabi…

---

06 **Hover value-frequency previews in Explorer** NEW 70

Adds hover previews for Top values, Rare values, and Count values that surface cached value frequencies from the current Explorer result without replacing the active pipeline, with an Open as query option for standalone analysis.

— Names the three preview types and the Open-as-query mechanism

Tenzir Platform v1.39.0: Explorer Usabi…

THINNER COVERAGE BELOW

---

07 **Temporal pipeline operators: session windows, lag, bounded reordering**

NEW 55

Adds three new pipeline processing primitives: session windows that group events by inactivity boundaries rather than fixed time intervals, a lag operator for accessing preceding-event values within a pipeline, and bounded event-time reordering to handle out-of-order events within a defined latency bound.

— Concepts named but no exact operator syntax or flags given product docs

---

08 **Statistical drift detection via mergeable models** NEW 45

Adds mergeable statistical models in TQL to build, store, merge, and compare data distributions, enabling drift detection across pipelines and time windows.

— Mechanism described in prose only, no operator name or syntax product docs

09 **Native Splunk HEC ingestion** NEW 35

Adds native Splunk HEC ingestion, enabling Tenzir nodes to receive data directly from Splunk HTTP Event Collector sources.

— Names the integration but no operator, config, or example given product docs

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-08-28 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260828.0 replaces external map tile servers with locally-rendered PMTiles maps (with admin-uploadable custom tiles and a new config key), and adds human-in-the-loop approval for Ask Ariane tool calls, new vulnerability propagation inference rules, a Score field on key entities, and map widget export.

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

01 Local PMTiles map rendering with admin-uploaded tiles

95

BREAKING

Maps are now served from a backend endpoint `/maps/world.pmtiles` using HTTP Range requests, removing any external map server dependency (including for air-gapped deployments) and superseding the old `map_tile_server_dark` / `map_tile_server_light` parameters, which are no longer used. A new `app:map_bundled_file_path` config key (env var `APP__MAP_BUNDLED_FILE_PATH`, default `./static/maps/world.pmtiles`) overrides the bundled file location, and administrators can upload a custom `.pmtiles` file via `Settings > Parameters > Map configuration` (stored in S3/MinIO, applied immediately with no restart) to replace bundled tiles with higher-resolution or region-specific data. Rendering automatically adapts to the platform's dark/light theme, and the platform supports low-zoom PMTiles extracts (zoom 0–6, ~30–80 MB) generated via `pmtiles extract --maxzoom=6`.

Generate a compact planet extract suitable for OpenCTI (zoom 0–6, ~30–80 MB) without downloading the full 137 GB source file.

```
$ pmtiles extract https://build.protomaps.com/20260722.pmtiles world.pmtiles --maxzoom=6
```

Override the bundled PMTiles file path when using a custom Docker image or non-standard layout.

```ini
APP__MAP_BUNDLED_FILE_PATH=/data/maps/custom-world.pmtiles
```

Upload a custom high-resolution PMTiles file to replace the bundled map for all users.

📍 1. Go to Settings › Parameters › Map configuration. 2. Click Upload and select your .pmtiles file. 3. The custom file is applied immediately with no restart required.

Point OpenCTI at an air-gapped or self-hosted PMTiles map file so the platform renders maps without external network access.

```yaml
app:
  map_bundled_file_path: '/opt/opencti/static/maps/world.pmtiles'
```

— Names endpoint, config key, env var, UI path and CLI command `7.260828.0`

---

02 **Vulnerability propagation inference rules** NEW 60

Adds a new inference rule deriving a System has Vulnerability relationship from combined System uses Software and Software has Vulnerability relationships, plus a related propagation rule linking vulnerabilities to Infrastructure through used software.

— Names exact relationship logic but rules run automatically, nothing to configure `7.260828.0`

THINNER COVERAGE BELOW

03 **Score field on key entities** NEW 50

Adds a Score field to Threat Actor, Intrusion Set, Malware, Incident, and Event entities.

— Names field and affected entity types, no further mechanism given `7.260828.0`

---

04 **Human-in-the-loop approval in Ask Ariane** NEW 45

Ask Ariane now requires analyst confirmation before the chatbot executes tool calls, adding a human-in-the-loop approval step to its tool-use flow.

— Describes the mechanism but no UI or config detail given `7.260828.0`

---

05 **Map widget export from dashboards** NEW 30

Enables export of map widgets from dashboards.

— Single-line description with no format, scope or mechanism `7.260828.0`

└──▷ **BREAKING ON UPGRADE**

! The `map_tile_server_dark` and `map_tile_server_light` configuration parameters are no longer used — external raster tile servers (e.g., `klokantech/openmaptiles-server`) can be removed from your deployment.

WAS THIS USEFUL?  

Recorded Future

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

Get started with the Recorded Future API. Covers authentication, base URLs, API key setup, and how to integrate threat intelligence into your security workflows.

> Recorded Future launched AI Alert Filtering, which automatically triages references inside fired Alerts by relevance and can auto-dismiss low-value alerts before they reach analysts.

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

01 **AI Alert Filtering for Alerts** NEW 65

A new per-rule AI Alert Filtering feature, enabled from the Alerts page in the Portal, automatically classifies every reference inside a fired Alert into High Relevance and Low Relevance sections before an analyst opens it. Analysts can set a custom intent per rule to sharpen prioritization beyond the rule's default intent (e.g. scoping to a specific organization), optionally enable auto-dismiss so Alerts with no references meeting the relevance threshold are dismissed before reaching the queue, and each Alert now gets an AI-generated title and summary highlighting whether anything demands immediate attention. Irrelevant references can also be excluded from API payloads to reduce downstream integration noise. It is available at no additional cost to all customers and enabled by default for new customers.

Enable AI Alert Filtering on a specific alert rule to start automatically triaging references by relevance without changing the underlying rule.

> 📍 In the Portal, go to Alerts › select a classic or custom alert rule › enable AI Alert Filtering for that rule.

Sharpen AI relevance decisions for a rule scoped to a specific organization by setting a custom intent, without rebuilding the rule from scratch.

> 📍 In the Portal, go to Alerts › select the alert rule › open AI Alert Filtering settings › enter a custom intent (e.g. 'This rule is for ACME Bank, not ACME Center').

— UI path and behavior given, but no API/config names for setup `launch-20260826-1c22a12a`

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo

Feedback ⓘ 2 RELEASES · 2026-08-25 → 2026-08-27 BLOG ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo Pro shipped Sensei CSPM for cloud misconfiguration scanning and direct remediation across AWS, Azure, and GCP, plus a Vulnerability Explorer with EPSS/KEV enrichment, KEV-aware SLA capping, and a beta root-cause clustering engine.

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

01 Fix in Cloud remediation with IaC PR fallback NEW

88

Fix in Cloud applies remediation directly through the cloud provider API using a separate write credential, independent of the read-only scan credential, and requires a mandatory human-approval dialog showing the exact action, target resource, and required cloud permissions before execution. When the affected resource is managed as code, Sensei instead opens a pull request via the existing IaC remediation path rather than touching the live resource, and every applied change is recorded in a Cloud Remediations ledger with a per-row Revert option.

Apply and later revert a Fix in Cloud action (e.g., blocking public S3 access) with human approval and a full audit trail.

📍 In DefectDojo Pro, open a cloud posture finding for an S3 bucket, click Fix in Cloud, review the approval dialog (action: block public access, resource: the specific bucket, required permissions listed), approve, then navigate to Sensei › Cloud Remediations to see the ledger entry and use the per-row Revert button if needed.

— Approval workflow, credential separation, and revert path fully described

launch-20260827-8a253852

02 Vulnerability Explorer with EPSS/KEV enrichment NEW

85

The Vulnerability Explorer is a deduplicated, one-row-per-vulnerability-ID view surfacing EPSS Score, EPSS Percentile, Known Exploited, Ransomware Used, and KEV Date columns alongside a Total Findings count that links directly to the filtered findings list. A `Sync KEV/EPSS data` button pulls the latest CISA KEV and FIRST EPSS data on demand in addition to the automatic daily sync, and the view distinguishes unenriched vulnerabilities (blank column) from confirmed non-exploited ones (red ✗ under Known Exploited).

Quickly surface the highest-exploitation-probability CVEs across your entire portfolio for same-day triage.

> 📍 In DefectDojo Pro, open the Vulnerability Explorer, click the `EPSS Score` column header to sort descending, then click `Total Findings` on the top row to jump directly to the filtered findings list for that CVE.

Force an immediate refresh of KEV and EPSS data after CISA publishes a new catalog update, without waiting for the nightly sync.

> 📍 In DefectDojo Pro, open the Vulnerability Explorer and click the `Sync KEV/EPSS data` button at the top of the page.

— Named columns and sync control, but no API endpoint given `launch-20260825-d1ed3339`

---

03 **Sensei cloud account onboarding and scanning** NEW 80

DefectDojo Pro adds Cloud Connection onboarding under Sensei to connect AWS Organizations, Azure subscriptions, and GCP projects in bulk using a single shared read-only credential that auto-discovers and onboards all visible accounts at once. On-demand posture scanning from the Sensei hub returns deduplicated, prioritized findings tied to the specific cloud resource (e.g. S3 bucket, security group) and updates them in place on re-scan.

Onboard all AWS accounts visible to an Organizations management credential and run an initial posture scan from the Sensei hub.

> 📍 In DefectDojo Pro, go to Sensei › Cloud Connections › Add Connection, select AWS, enter a read-only Organizations credential, and choose `Discover and onboard all member accounts`. Then navigate to Sensei › Hub and select `Run Scan` against the newly onboarded accounts.

— Clear UI path and mechanism, but no API/CLI surface named `launch-20260827-8a253852`

04 **Root Cause Correlation (Beta) clustering** NEW 80

Root Cause Correlation (Beta) groups findings that share the same component-at-version, CVE, infrastructure resource, or weakness class at the same URL into ranked clusters showing active findings, affected assets, and aggregate risk. It is enabled by a superuser under Settings > Feature Flags.

Enable Root Cause Correlation (Beta) so that clusters of findings sharing the same vulnerable component are grouped and prioritized together.

> 📍 In DefectDojo Pro, navigate to Settings › Feature Flags and enable Root Cause Correlation (requires superuser account).

— Clustering logic and exact enable path both specified `launch-20260825-d1ed3339`

THINNER COVERAGE BELOW

05 **Cap KEV-listed SLAs to CISA due date** NEW 55

SLA configuration supports Cap by KEV due date, which tightens a KEV-listed finding's remediation deadline to CISA's official due date for that CVE.

— Config option named but no navigation path or example given `launch-20260825-d1ed3339`

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ 4 RELEASES · 2026-08-26 → 2026-08-28 BLOG ↗

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale's biggest window shipment is a Privileged Access Management (PAM) beta for brokering access to databases, RDP, SSH, Kubernetes and HTTPS services, alongside Aperture reaching GA as an AI agent/LLM gateway with new MCP endpoints; it also added Control D-powered DNS filtering and a container startup timeout knob.

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

01 Aperture AI gateway for agents and LLM sessions NEW

88

Tailscale introduced Aperture, a gateway for securing and managing AI agents and LLM sessions, with cost controls, request/response hooks, guardrails, logging, MCP support, and API proxying. It exposes `Tailscale` and `Tailscale SSH` MCP endpoints, letting coding agents add tailnet nodes or access them via Tailscale SSH after user approval, without manually distributing keys. `Projects` group chats with shared instructions, tool access, and tailnet nodes to reduce per-chat re-prompting, and default tool permissions can be defined across all Aperture chats to control baseline tool access. New Aperture instances include initial tokens for major open-weight and closed models, with additional model tokens purchasable directly inside Aperture without supplying third-party API keys.

When you want an AI agent to deploy a service to a tailnet node and configure SSH access without manually copying keys, use the new Tailscale SSH MCP endpoint through Aperture.

> 📍 In Aperture, go to MCP Endpoints, enable the `Tailscale SSH` MCP, then open a chat and instruct your agent to access a named tailnet node via Tailscale SSH — Aperture will prompt you to approve any new machine before the agent proceeds.

When running recurring project work (e.g. homelab automation), create a Project in Aperture to share instructions, tool access, and tailnet nodes across all related chats so you stop re-prompting context each session.

> 📍 In Aperture, go to Projects › New Project, set shared initial instructions and select the tailnet nodes and tools this project's chats may access, then open chats inside the project to inherit those defaults.

— Named MCP endpoints and exact UI navigation steps provided `launch-20260826-f7b3ed6e` `changelog-20260828-f4636e7a`

02 DNS filtering via Control D add-on NEW

86

Tailscale added DNS filtering via Control D as a purchasable add-on, applied per-user, per-group (`group:`), per-tag, or per-device through the existing ACL file using the `nodeAttrs` attribute (e.g. `controld:<rule-id>`). Covered devices route DNS through Control D over encrypted DNS, enforcing rulesets for malware, phishing, content categories, and custom block/allow/redirect rules; rules remain manageable via the Control D Dashboard or API independently of Tailscale.

Apply a Control D DNS filtering rule to all employees in your tailnet to block malicious and phishing domains.

```json
{
  "nodeAttrs": [
    {
      "target": ["group:employees"],
      "attr": [
        "controld:o13v9z7zns"
      ]
    }
  ]
}
```

Add Control D as a nameserver to your tailnet via the Tailscale admin console to route DNS through Control D's filtering infrastructure.

> 📍 In the Tailscale admin console, go to DNS › Nameservers and add Control D as a nameserver. Then open your Control D dashboard to assign or create a filtering rule.

— Named ACL attribute with a runnable JSON config example `launch-20260828-a7200f1e`

---

03 **TS_BOOT_TIMEOUT container startup env var** NEW 85

New `TS_BOOT_TIMEOUT` environment variable controls how long the Tailscale container waits for its initial connection to the coordination server, accepting durations like `90s` or `3m` and defaulting to `60s`.

Extend container startup grace period in slow or heavily loaded environments where the coordination server connection takes longer than the default 60 seconds.

```
$ TS_BOOT_TIMEOUT=3m docker run --rm -e TS_BOOT_TIMEOUT=3m
tailscale/tailscale
```

Set `TS_BOOT_TIMEOUT` in a Kubernetes deployment so the sidecar container waits 90 seconds before failing on slow cluster network initialization.

```yaml
env:
  - name: TS_BOOT_TIMEOUT
    value: "90s"
```

— Exact env var, default, and runnable docker/k8s examples `changelog-20260828-f4636e7a`

---

04 **Privileged Access Management (PAM) beta** NEW 83

Tailscale introduced Privileged Access Management (PAM) in beta, managed from the Tailscale admin console, supporting database, RDP, HTTPS, SSH, and Kubernetes Service types with session recording, playback, and audit logs of user access. Policies for user access and in-service permissions are set through grants, with secrets store integration for credential injection, and PAM connectors can be deployed inside private environments without installing agents on each resource. Just-in-time access requests and approvals are supported via Slack, scoped to a specific resource and time window, and per-service access policies control users, groups, time windows, and permissions at finer granularity than subnet-level network access. A browser client allows access to PAM services without requiring a Tailscale client download, custom domain management is available for public HTTP PAM services, a custom Amazon S3 bucket can be configured for PAM recording storage, and service accounts can be created for machine identities to interact with the PAM API.

— Extensive mechanism and scope but no exact CLI/API commands given

changelog-20260828-f4636e7a launch-20260827-8403c4f5

THINNER COVERAGE BELOW

05 **Peer relay health checks on AWS fixed** IMPROVED 55

Peer relays on AWS now use HTTP instead of TCP for pod health checks, preventing all targets from being reported as unhealthy, and are reachable regardless of the availability zone their pod is scheduled in; pinning Elastic IPs requires pinning pods to the same zone using a Network Load Balancer.

— Clear before/after fix but no reader-facing action changelog-20260828-f4636e7a

WAS THIS USEFUL? 👍 👎

Zscaler

Feedback

1 RELEASE · 2026-08-27

BLOG ↗

Zscaler is a cloud-based security platform that inspects and filters internet traffic to protect users from cyber threats and data loss.

> Zscaler's Zero Trust Browser gains WebMCP Security Controls, a new pipeline to discover, monitor, and enforce policy on AI agent tool calls running inside web pages.

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

01 **WebMCP Policy Enforcement** NEW 95

Inline policies match on any of the seven `webmcp.*` attributes and apply one of three effects — Allow, Monitor, or Block — intercepting tool calls before the site's handler runs. Rules can inspect tool identity (`webmcp.tool.name`), argument content (`webmcp.tool.arguments`), provider origin (`webmcp.tool.provider_origin`), and cross-origin or untrusted-content signals (`webmcp.tool.is_cross_origin`, `webmcp.tool.has_untrusted_content_hint`) for DLP-style argument-level inspection; blocked calls return a structured failure to the agent with the exact rule reason (e.g., `webmcp.tool.name is equal to UPDATE_CART`) logged as a Low-severity Block detection.

Block a specific AI agent capability by tool name — for example, prevent any agent from invoking `update_cart` on any covered site, with the block surfaced inline to the user.

📍 In the Zscaler admin console, create a policy named Block WebMCP with condition: `webmcp.tool.name` is equal to update_cart, effect: Block. The Detections view will show a Low-severity Block entry with Reason: `webmcp.tool.name` is equal to `UPDATE_CART` on every blocked invocation.

Apply argument-level (DLP-style) inspection to block agent tool calls whose JSON payload matches a keyword, while allowing the same tool with clean payloads.

📍 In the Zscaler admin console, create a policy rule with condition: `webmcp.tool.arguments` contains book, effect: Block. Calls to update_cart with payloads such as {"cart":{"line_items":[{"handle":"book-cat-graphic-tee-apple-blossom-26f12103","quantity":1}]}} are blocked and the full argument payload is captured verbatim in the detection record.

— Includes exact admin-console rule syntax and observed block behavior

2026-08-27-3c00b999

02 **WebMCP Runtime Monitoring** NEW 78

Every AI agent tool invocation is captured with full context — the seven `webmcp.*` attributes plus user identity, device posture, extension version, browser, IP, geolocation, and URL — and surfaced in the Zscaler admin console as a searchable, filterable, exportable detection in the same event stream as existing platform telemetry.

Start in monitor mode to baseline all AI agent tool calls across sites before writing any block rules — every `webmcp.*` attribute is logged as an Informational/Allow detection without disrupting users.

> 📍 In the Zscaler admin console, create a policy named `Monitor WebMCP` with all seven webmcp.* attributes enabled, default effect set to Allow, and assign it to the target user group. Open the Detections view to stream every WebMCP tool registration and invocation with full user, device, and argument context.

— Includes exact admin console steps to baseline traffic in monitor mode

2026-08-27-3c00b999

---

03 **WebMCP Capability Discovery** NEW 70

The Zero Trust Browser Extension enumerates every tool a WebMCP-enabled site registers before any can run, capturing seven attributes per tool — `webmcp.tool.name`, `webmcp.tool.description`, `webmcp.tool.arguments`, `webmcp.tool_count`, `webmcp.tool.provider_origin`, `webmcp.tool.is_cross_origin`, and `webmcp.tool.has_untrusted_content_hint` — building a structured inventory of the agentic surface per site and user.

— Names all seven attributes but no direct configuration step shown 2026-08-27-3c00b999

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Immuta

Feedback ⓘ 1 RELEASE · SEEN 2026-08-28 NOTES ↗

Immuta is a data governance and access control platform that enables secure, policy-driven data sharing while maintaining compliance and privacy.

> Immuta introduces Intents, a new context-aware access control model now in public preview for Databricks Unity Catalog, alongside a new Connections API and a wave of connection-management guides for BigQuery, Starburst (Trino), Snowflake, and AWS Lake Formation; subscription policies now cover Unity Catalog metric views, and lockout policies now trigger on policies referencing nonexistent columns.

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

01 **Intents for context-aware access control** NEW 73

Intents let data consumers assume a declared reason (intent) to unlock policy-conditional access to sensitive fields such as PII, replacing birthright attributes with intent-specific attributes while the intent is active. Intents can be created, edited, and deleted programmatically via the Immuta v1 API (documented in the Manage Intents API reference guide) or managed in the UI, and are now in public preview for Databricks Unity Catalog.

— Names API, guide and Unity Catalog scope but no exact endpoint syntax snapshot-20260828

02 **Connection management guides for BigQuery, Starburst, Snowflake and tags**

NEW 70

New how-to guides cover registering, removing, and managing connections across integrations: a Google BigQuery viewless integration for registering and governing BigQuery data without view-based policy enforcement (with guides on policy enforcement, data access, and security and compliance); a Starburst (Trino) integration section covering registering, removing, and managing Trino connections, including legacy connection support and system API key rotation; new Snowflake connection how-to guides for registering, removing, and managing connection settings, plus reference pages for Snowflake Low Row Access Policy Mode and Snowflake Table Grants; and a Manage Connection Tags how-to guide for tagging registered connections.

— Names every integration and guide but gives no commands or exact steps product docs

THINNER COVERAGE BELOW

03 **Connections API and AWS Lake Formation registration guide**

NEW 57

A new Connections API allows programmatic registration and management of data connections, documented in the Connections API reference guide. A dedicated AWS Lake Formation Connection Registration Payloads Reference Guide documents the parameters and payloads needed to register an AWS Lake Formation connection via this API.

– Names the API and reference guide but no method/path or sample payload `product docs`

---

04 **Lockout policies now apply to missing column references**

50

BREAKING

Lockout policies now apply to data sources governed by a policy that references a nonexistent column or column tag; previously no lockout was applied in those cases.

– Clear before/after behavior but no remediation steps given `snapshot-20260828`

---

05 **Subscription policies for Unity Catalog metric views** IMPROVED

48

Subscription policy support now extends to Databricks Unity Catalog metric views, allowing access provisioning and governance for that object type through existing Immuta subscription policies.

– Describes the change but not how to configure it `snapshot-20260828`

└──▷ **BREAKING ON UPGRADE**

! Lockout policies will now apply to data sources governed by a policy that references a nonexistent column or column tag; previously no lockout was applied in those cases.

WAS THIS USEFUL?  

Microsoft Purview

Feedback

1 RELEASE · 2026-08-27

BLOG ↗

Microsoft Purview is a unified data governance and compliance platform that discovers, maps, and manages data across organizations.

> Microsoft Purview raised the daily throughput cap for auto-labeling policies on SharePoint and OneDrive files fivefold.

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

01 **Higher auto-labeling capacity for SharePoint and OneDrive**

IMPROVED 55

Auto-labeling policies now support processing up to 500,000 SharePoint and OneDrive files per day, up from a previous limit of 100,000, expanding sensitivity-label coverage at scale.

— Names exact before/after limits but no config or rollout detail. 2026-08-27-fbcac0d9

WAS THIS USEFUL? 👍 👎

Varonis

Feedback

2 RELEASES · 2026-08-03 → 2026-08-07

BLOG ↗

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis Atlas gained two AI-security capabilities this window: a real-time inline enforcement integration with Claude Inference Hooks, and a new Agent Intent-Based Access Control system that detects and quarantines AI agents drifting from their instructions.

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

01 **Agent Intent-Based Access Control (IBAC)** NEW 84

New Agent IBAC capability in Varonis Atlas compares an AI agent's received instruction against its reasoning, tool calls, and data access in real time, alerting, blocking, modifying, logging, or routing actions to a human for approval when drift is detected. Drift detection runs at three tunable sensitivity levels — lenient, balanced (default), and strict — and full-session evaluation tracks every prompt, response, and tool call across a conversation (keyed by conversation ID) to catch gradual or multi-turn drift, including jailbreaks spread across individually benign prompts. Teams can write session policies in plain language and set how many events must accumulate before evaluation runs; quarantine blocks the violating identity for a customer-configured window from minutes to a full day, with an admin console listing quarantined identities and controls to lift, extend, or make quarantine permanent. A full audit trail covers every prompt, model response, tool execution, intent assessment, and guardrail action, viewable via a conversation view and an expanded execution view of underlying tool calls.

— Named sensitivity levels, quarantine windows, and admin console give a clear starting point

launch-20260803-7ef121fc

02 **Claude Inference Hooks integration in Atlas** NEW 69

Varonis Atlas now integrates with Claude Inference Hooks so it sits in the request path, receiving the full conversation transcript from Anthropic and returning an allow-or-deny verdict before inference runs — denied prompts never reach Claude. It inspects prompt content and attachments in real time against customizable policies covering PII exposure, malicious URL detection, and prompt injection, with enforcement extended across Claude Chat, Claude Design, Claude Enterprise, Claude Platform, Claude Cowork, and Claude Code, plus a complete audit trail of every prompt, response, and tool execution alongside the actions Atlas took.

— Clear mechanism and scope named, but no config keys or API details given

launch-20260807-fb69b76b

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Material Security

Feedback ⓘ 1 RELEASE · SEEN 2026-08-28

BLOG ↗

Material Security is an email security platform that detects and prevents advanced threats, data loss, and account compromise in cloud email environments.

> Material Security launched an MCP (Model Context Protocol) server that lets AI agents triage email security issues through the same permissions, rate limits, and audit trail as human users.

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

01 MCP server for AI-agent email triage NEW

95

Material added an MCP server, enabled via an `'MCP Server'` toggle in Settings (default off), exposing four tools to AI agents: `list_issues`, `get_issue`, `patch_issue`, and `get_message`. These are backed by new API v1 endpoints with guaranteed schema parity to the REST API, so anything callable via MCP is also callable via REST. Agents authenticate through OAuth (rolled out in v1.60) using Material's existing SAML/Google/Microsoft login instead of long-lived API tokens, with supported harnesses including Cursor, Claude, ChatGPT, Codex, and Grok. Tool-level permissions are enforced per user and re-checked at execution time — `patch_issue` is hidden without write access and `get_message` is gated by message-read permissions — and the server is gated at the tenant level, available only when every tenant a user belongs to has opted in. Requests are rate-limited per tenant via Redis with fail-fast behavior so a runaway agent gets an immediate error instead of queuing silently, and every MCP tool call (tool name, parameters, caller identity, success/error, and for `patch_issue` the status, classification, and remediation changes) is written to Material's audit log alongside regular API and UI activity.

Enable the MCP server for your tenant so AI agents can connect to Material's email security tools.

> 📍 In the Material UI, go to Settings › MCP Server and toggle MCP Server on, then copy the one-click connection snippet for your AI agent harness (e.g. Cursor, Claude, or ChatGPT).

— Names every tool, endpoint parity, toggle, and safeguard with setup steps.

launch-20260828-d16fc689

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ 1 RELEASE · 2026-08-28 BLOG ↗

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0 has taken Custom Token Exchange to General Availability, letting applications swap an existing token for Auth0-issued tokens under custom validation logic, with full Dashboard management and delegated-authorization support.

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

01 **Custom Token Exchange reaches General Availability** NEW

74

Custom Token Exchange lets applications exchange an existing token — issued by Auth0 or an external identity provider — for Auth0 access, ID, and refresh tokens, validated by a Custom Token Exchange Action. It supports delegated authorization so a service, AI agent, or support agent can act on behalf of a user, including calling APIs or establishing a web session via Session Delegation. The Auth0 Dashboard now supports creating and managing Custom Token Exchange Profiles and their associated Actions directly, and a shared transaction context lets you set metadata in a Custom Token Exchange Action and read it from a Post-Login Action to customize resulting tokens based on subject or actor token information. It is available to all Enterprise, B2B Professional, and B2C Professional customers, and included in the Trial period for free tenants.

.
.
.
.

— Mechanism and named surfaces given, but no code sample or exact config steps.

Custom Token Exchange is now Generally …

WAS THIS USEFUL? 👍 👎

Token Security

Feedback

2 RELEASES · 2026-08-26 → 2026-08-27

BLOG ↗

Token Security is a blockchain security analysis platform that audits smart contracts and tokens for vulnerabilities, scams, and security risks.

> Token Security shipped two AI-native additions this window: an MCP server exposing NHI inventory and remediation to any MCP-compatible AI client, and Enzo, a natural-language builder for live identity security applications.

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

01 **NHI MCP Server for AI clients** NEW 80

A new Token MCP Server exposes Token's NHI inventory, findings, risk data, and remediation guidance to any MCP-compatible AI client (Claude, ChatGPT, Gemini, Cursor) via the Model Context Protocol standard. It supports natural-language queries across NHI inventory, permission structures, authentication methods, environment criticality, ownership, and usage patterns, returning prioritized risk views, and can generate customer-specific de-provisioning scripts and CLI commands on demand. It also enables cross-cloud identity analysis (e.g. GCP service accounts consumed by AWS resources) and connects into agentic workflows where AI agents autonomously retrieve findings, generate JIRA tickets with pre-filled scripts, and send Slack notifications — accessible from both the Token Portal and external AI chat/agent apps.

— Names specific AI clients, protocol, cross-cloud example, and integrations, but no exact commands

launch-20260827-096a8588

02 **Enzo AI-native application builder** NEW 72

Enzo is a new AI-native application builder inside the Token Security platform that generates live, runnable identity security applications from natural-language prompts, with no code required. Generated apps connect to live identity data across human identities, non-human identities, AI agents, SaaS, cloud infrastructure, secrets, and credentials, and can be interactive dashboards, approval flows, remediation tools, or governance apps rather than static reports, with automation actions including Slack notifications, email workflows, ticketing, approvals, and lifecycle management. Apps run in sandboxed, tenant-isolated environments with per-session credential scoping and full audit logging, and can be persisted, refined through further prompts, and shared across teams.

— Describes mechanism and scope in detail but gives no exact command or path

launch-20260826-e4806597

WAS THIS USEFUL?  

Didit

Feedback

1 RELEASE · SEEN 2026-08-28

NOTES ↗

> Didit's biggest shift this window is Networks, which clusters applicants, businesses, and transaction parties across an entire organization's applications into single fraud-ring investigations; alongside it, Didit shipped a full transaction-monitoring rules Management API (with MCP tools and AI-assisted rule migration), a new AML screening endpoint, a wallet-screening PDF export, and a standalone document OCR/fraud-check endpoint.

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

01 Networks: cross-application fraud-ring clustering NEW

95

Networks automatically clusters applicants, businesses, and transaction parties across every application in an organization that share device, IP, biometric, document, or contact signals, surfacing a fraud ring as one investigation instead of dozens of unrelated flags. High-degree hub signals (shared office IP, kiosk device, data-center IP) are flagged for analyst review rather than auto-declined, distinguishing legitimate high-traffic environments from fraud. Analysts can create a case anchored to a network for handoff to compliance, or blocklist a member's identity or an individual signal directly from the network detail page so future matching sessions auto-decline. Every signal value except selfie-background and face-similarity is stored as a keyed hash of its normalized value, and clustering runs on production data only — sandbox applications and sessions never join a network. Exposed via API through `GET /v3/organization/{organization_id}/application/{application_id}/networks/`, `.../networks/{network_uuid}/`, `.../networks/{network_uuid}/graph/`, `.../networks/{network_uuid}/map/`, `.../networks/{network_uuid}/members/`, `.../networks/{network_uuid}/signals/`, `.../networks/{network_uuid}/timeline/`, and membership lookups at `.../business-sessions/{business_session_id}/networks/`, `.../businesses/{vendor_business_uuid}/networks/`, `.../sessions/{session_id}/networks/`, `.../transactions/{transaction_uuid}/networks/`, `.../transactions/{transaction_uuid}/parties/{transaction_party_uuid}/networks/`, and `.../users/{vendor_user_uuid}/networks/`.

Investigate a flagged network in the console by narrowing the graph to the immediate connections of one suspicious applicant.

> 📍 In the Didit console, open Networks › select the network › Graph tab, then click an applicant node and set the hop limit to 1, 2, or 3 to isolate their direct

connections.

— Full mechanism, thirteen named endpoints, and a console workflow given. `3.0.0`

02 **AML screening endpoint** NEW 95

New `POST /v3/aml/` endpoint screens a person or company against PEP, sanctions, watchlist, and adverse-media datasets, returning `aml.status` (`Approved` / `In Review` / `Declined`) and per-hit `match_score` and `risk_score` fields. The two-score model is configurable per-request via `aml_match_score_threshold` (default 93), `aml_score_review_threshold` (default 100), `aml_score_approve_threshold` (default 80), and `aml_name_weight` / `aml_dob_weight` / `aml_country_weight` (must sum to 100, defaults 60/25/15). `include_adverse_media=true` adds adverse-media enrichment (~5–25s latency; a 60s+ client timeout is recommended), and `include_ongoing_monitoring=true` keeps re-screening a profile over time (requires `save_api_request=true`, enforced with a `400` otherwise). `save_api_request` (default `true`) persists screenings as sessions retrievable via `GET /v3/session/{sessionId}/decision/` and fires `status.updated` webhooks. The endpoint also gained an `Idempotency-Key` header parameter to prevent duplicate charges on retry.

Screen a person against PEP and sanctions lists with retry safety — use `Idempotency-Key` so a timeout does not trigger a duplicate AML charge.

```
$ curl -X POST https://verification.didit.me/v3/aml/ \
 -H 'x-api-key: YOUR_API_KEY' \
 -H 'Content-Type: application/json' \
 -H 'Idempotency-Key: aml-user-1234-2026-08-20' \
 -d '{
  "full_name": "Antonio Ejemplo Modelo",
  "entity_type": "person",
  "date_of_birth": "1972-02-29",
  "nationality": "ES",
  "include_adverse_media": false,
  "vendor_data": "user-1234"
 }'
```

— Full parameter set, defaults, and a runnable curl example. `3.0.0`

03 **Transaction-monitoring rules Management API** NEW 90

Adds `/v3/transactions/rules/` endpoints to create (`POST /v3/transactions/rules/`), list (`GET /v3/transactions/rules/`), retrieve (`GET /v3/transactions/rules/{rule_uuid}/`), update (`PATCH /v3/transactions/rules/{rule_uuid}/`), delete (`DELETE /v3/transactions/rules/{rule_uuid}/`), and backtest (`POST`

`/v3/transactions/rules/backtest/`) transaction-monitoring rules programmatically, complementing the Business Console. A rule library can be browsed (`GET /v3/transactions/rules/library/`), and presets installed or removed via `POST /v3/transactions/rules/install/` and `DELETE /v3/transactions/rules/install/`. The list endpoint is filterable by `source` (`PRESET` / `CUSTOM`), `category`, `mode` (`ACTIVE` / `DISABLED` / `TEST`), and `bundle`, with `ordering` support on `run_count`, `approved_pct`, `reviewed_pct`, `declined_pct`, `latest_triggered_at`, `title`, `created_at`, and `source` (each with an optional `-` prefix for descending).

List only active custom transaction-monitoring rules, ordered by most recently triggered, to review which rules are firing in production.

```
$ curl -G 'https://verification.didit.me/v3/transactions/rules/'
\
  -H 'x-api-key: YOUR_API_KEY' \
  --data-urlencode 'source=CUSTOM' \
  --data-urlencode 'mode=ACTIVE' \
  --data-urlencode 'ordering=-latest_triggered_at'
```

— Endpoints, filters, and a runnable curl example fully specified. `3.0.0` `snapshot-20260828`

04 **Standalone document OCR and fraud-check endpoint** NEW

78

New endpoint accepts `front_image` and optional `back_image`, automatically classifies 14,000+ document types from 220+ countries, and returns extracted holder fields, parsed MRZ, an `Approved` / `Declined` `status`, and a `warnings` list. It exposes per-request risk-action fields `invalid_mrz_action`, `inconsistent_data_action`, and `expiration_date_not_detected_action` (each `DECLINE` or `NO_ACTION`), and `perform_document_liveness=true` screens for `SCREEN_CAPTURE_DETECTED`, `PRINTED_COPY_DETECTED`, and `PORTRAIT_MANIPULATION_DETECTED`. Its `save_api_request` field (default `true`) persists the call as an API-type session retrievable with `GET /v3/session/{sessionId}/decision/`.

— Rich field detail but no runnable example provided. `product docs`

05 **Wallet-screening PDF export** NEW

75

New `POST /v3/wallet-screening/pdf/` endpoint re-renders a prior wallet-screening result as a Didit-branded PDF by reposting the exact JSON body (including `report_signature`) returned by `POST /v3/wallet-screening/`, with no additional provider call or billing. The underlying `POST /v3/wallet-screening/` endpoint's response schema was also changed in the same release, though the specific fields altered are not detailed.

Export an already-run wallet screening as a signed PDF for audit — repost the unmodified JSON body (with `report_signature`) returned by `POST /v3/wallet-screening/`.

```
$ # $RESULT is the saved JSON body from POST /v3/wallet-
screening/
curl -X POST 'https://verification.didit.me/v3/wallet-
screening/pdf/' \
  -H 'x-api-key: YOUR_API_KEY' \
  -H 'Content-Type: application/json' \
  -d "$RESULT" \
  -o wallet_screening_report.pdf
```

— Endpoint mechanism and a runnable export example given. 3.0.0

THINNER COVERAGE BELOW

06 **AI-assisted transaction-rule migration** NEW 50

The Didit assistant, or any MCP-connected agent, can ingest a rule set exported from another provider as CSV or Excel, map it to Didit's rule schema, backtest it against the organization's transaction history, and create the rules in test mode for review before activation.

— Workflow described but no concrete command or endpoint given. snapshot-20260828

07 **MCP tools for transaction rules** NEW 45

Introduces nine `didit_transaction_rule_*` MCP tools covering list, create, update, delete, backtest, and library workflows for transaction-monitoring rules, usable by any MCP-connected AI agent.

— Tool categories named but no exact tool names or commands given. snapshot-20260828

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-28 BLOG ↗

Clerk is an authentication and user management platform that handles sign-up, login, and identity verification for applications.

> Clerk now lets developers customize the reverification window that governs how recently a user must have authenticated before performing sensitive actions.

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

01 **Customizable reverification window** NEW 31

Clerk adds the ability to customize the reverification window, defining how recently a user must have authenticated before sensitive actions are permitted.

— No config key, API, or default value given. Customize the reverification window (20…

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ 1 RELEASE · 2026-08-26 NOTES ↗

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 v3.21.3 sharpens audit gap-finding and extends the workflow engine's reach over objects, while adding a new AI governance framework to its library.

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

01 **Audit tree filter for uncovered requirements** NEW 63

The audit tree now supports filtering to show only requirements without controls or evidences, letting auditors isolate compliance gaps quickly by enabling Requirements without controls or evidences in the filter panel.

Quickly surface audit gaps by filtering the audit tree to show only requirements that have no controls or evidences mapped.

📍 In an open audit, navigate to the Audit Tree view, then use the filter panel to enable Requirements without controls or evidences to isolate uncovered items.

— Names the exact filter and navigation path v3.21.3

THINNER COVERAGE BELOW

02 **Workflow engine object registry and triggers** IMPROVED 56

The workflow engine's `read_objects` registry is expanded to include requirements, risk objects, and validations, letting workflows act on a broader set of objects. The engine also gains the ability to write objects, schedule actions by date, and detect bulk edits as triggers.

— Names registry key and new capabilities but no usage steps v3.21.3

03 **UNESCO AI Maturity Framework library entry** NEW 25

Adds the UNESCO AI Maturity Framework as a new library entry, supporting AI governance assessments.

— Bare name with no mechanism or usage detail v3.21.3

WAS THIS USEFUL? 👍 👎

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ 2 RELEASES · 2026-08-28 BLOG ↗

RegScale is a compliance automation and GRC platform that helps organizations manage regulatory requirements, assessments, and security governance.

> RegScale's CLI gained new FedRAMP SSP import commands while the core platform introduced a Key Risk Indicator monitoring system and a consolidated dashboard experience.

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

01 **FedRAMP SSP import commands in CLI** NEW 60

CLI 6.42.32 adds a new FedRAMP SSP `interconnect` import command and a new FedRAMP SSP `cryptographic module` import command, and adds an optional SSP ID parameter to the FedRAMP `docx` import command so an existing System Security Plan is updated in place rather than duplicated.

— Names three specific CLI import commands but no full syntax given

CLI 6.42.32 (2026-08-28)

THINNER COVERAGE BELOW

02 **Key Risk Indicator (KRI) monitoring system** NEW 30

Introduces a Key Risk Indicator (KRI) system for organizations to monitor risk metrics.

— States purpose only, no metrics, thresholds or UI path given

RegScale 6.33.3.0 (2026-08-28)

03 **Consolidated dashboard experience** IMPROVED 20

Consolidates the dashboard experience into a unified view for accessing and reviewing organizational metrics.

— Generic description with no specifics on what changed or how to access it

RegScale 6.33.3.0 (2026-08-28)

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

Scrut Automation

Feedback ⓘ 1 RELEASE · SEEN 2026-08-28

BLOG ↗

> Scrut launched Shadow AI Governance, a new module that discovers AI tool usage across an organization, classifies it, and lets teams set enforceable governance statuses with audit trails.

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

01 **AI application discovery and classification** NEW 68

A new Shadow AI Governance module discovers AI applications employees are using via SSO integrations (Google Workspace, Microsoft Entra) and desktop sources (supported MDM integrations or the Scrut Agent). Discovered tools are automatically classified as AI Apps (tools where employees input organizational data) or AI Builders (tools that connect to technical environments such as repositories and APIs), each surfacing distinct review questions.

— Names discovery sources and classification categories but no UI/API path

launch-20260828-a5da0428

02 **Governance status and restriction remediation** NEW 65

Each discovered application can be assigned an explicit governance status of Managed, Ignored, or Restricted, with Restricted decisions requiring a recorded justification to build an auditable trail. When an app is Restricted, Scrut identifies affected employees and notifies them with removal instructions, including third-party access revocation for SSO-discovered apps and uninstall guidance for desktop apps.

— Concrete statuses and remediation steps, but no exact config or API

launch-20260828-a5da0428

THINNER COVERAGE BELOW

03 **Shared AI footprint view** NEW 33

Provides a shared AI footprint view for security, IT, and GRC teams, separate from an approved software inventory, showing which employees are associated with each discovered application.

— Describes the view's purpose but no navigation or mechanism detail

launch-20260828-a5da0428

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-27 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

> OpenShell v0.0.115 adds OTLP-based distributed tracing with gateway identity tagging, a native Windows MXC compute driver, OCI supply-chain attestations, and a streamlined local Kubernetes gateway workflow.

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

01 **OTLP trace export with gateway identity** NEW 58

OpenShell now exports Kubernetes driver traces over OTLP, letting distributed tracing pipelines ingest gateway activity. Exported traces are tagged with gateway identity so multi-gateway deployments can filter and correlate traces per gateway.

— Names OTLP and gateway tagging but no config or endpoint given. v0.0.115

02 **Native Windows MXC compute driver** NEW 45

Adds a native Windows MXC compute driver with full server wiring, extending OpenShell's supported compute platforms to Windows MXC.

— Names the platform but no setup steps or flags. v0.0.115

03 **OCI SBOM and provenance attestations** NEW 40

Publishes OCI SBOM and provenance attestations alongside release artifacts for supply-chain verification.

— Names artifact types but no retrieval command shown. v0.0.115

04 **Unified local Kubernetes gateway workflow** IMPROVED 23

Unifies the local Kubernetes gateway workflow, streamlining Kubernetes-backed development environments.

— Vague description with no concrete mechanism or command. v0.0.115

WAS THIS USEFUL? 👍 👎

ToolHive

Feedback ⓘ 2 RELEASES · 2026-08-26 → 2026-08-27 NOTES ↗

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

> ToolHive's two releases centered on plugin supply-chain security — Sigstore-verified plugin artifacts and signer-rotation blocking on upgrades — alongside a CVE fix closing a CORS hole in the `thv serve` management API, new private-CA trust for OIDC/OAuth2 upstreams, and a batch of smaller hardening and bug fixes spanning skill sync, Virtual MCP, and the workloads API.

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

01 **Signer-rotation blocking for plugin upgrades** NEW 95

`thv ai-plugin upgrade` now blocks when the signature identity of an upgrade differs from the one recorded in the project lock file (or is unsigned), exiting with code 4 and reporting `signer-change-blocked`. The new `--allow-signer-change` flag re-records the new signer identity to approve the change; the check is controlled by `TOOLHIVE_PLUGINS_LOCK_ENABLED`, and lock entries with no recorded provenance are unaffected.

Upgrade a plugin whose signing identity has legitimately rotated, explicitly approving the new signer so it is recorded in the lock file.

```
$ thv ai-plugin upgrade --allow-signer-change <plugin-name>
```

— Names flag, exit code, message, env var, and gives runnable command v0.46.0

02 **CORS hardening for thv serve management API** BREAKING 95

The `thv serve` management API for TCP listeners now enforces `Content-Type: application/json` on state-changing requests with a body, rejecting non-conformant calls with `415 Unsupported Media Type`, and validates `Origin` against a loopback-only allowlist (non-loopback binds log a `WARN` and pass through), blocking CORS-based cross-origin attacks tracked as GHSA-xv9h-79wp-q9w6. UNIX-socket paths (`--socket=<path>`) are unaffected.

Call the `thv serve` management API correctly after v0.45.0 requires `Content-Type: application/json` on mutating TCP requests.

```
$ curl -X POST http://127.0.0.1:8080/api/v1beta/workloads \
  -H 'Content-Type: application/json' \
  -d '{"name":"fetch","image":"ghcr.io/example/fetch:latest"}'
```

— Names header, status code, allowlist behaviour, CVE, and a curl example `v0.45.0`

---

03 **Private CA trust for OIDC/OAuth2 upstreams** NEW 85

A new `caBundleRef` field on OIDC and OAuth2 upstream specs lets the embedded auth server trust a private CA for discovery, token, user-info, and dynamic client registration calls to an in-cluster identity provider. Note: the `operator-crds` chart must be updated to 0.46.0 before or with the operator chart, since a stale CRD silently prunes `caBundleRef` from applied resources instead of rejecting it.

Trust an internal CA for an in-cluster OIDC provider so the embedded auth server can reach it for discovery and token calls.

```yaml
caBundleRef: my-internal-ca-secret
```

— Names field and every call path it covers, plus upgrade caveat `v0.46.0`

---

04 **Keyless signing required for skill push** BREAKING 80

`thv skill push` now signs artifacts keylessly by default and requires exactly one of `--key`, `--identity-token`, or `--no-sign`; specifying both `--key` and `--no-sign` now returns `400`.

Push a skill artifact with keyless signing using `--identity-token` now that `thv skill push` requires exactly one signing option.

```
$ thv skill push --identity-token "$(cat /var/run/secrets/token)"
./my-skill
```

— Names all three flags and the error, with a runnable example `v0.45.0`

---

05 **Qoder client and skill sync scope change** BREAKING 80

Adds `qoder` as the 18th skill-supporting client; `thv skill sync` run without `--clients` now targets every skill-supporting client, including `.qoder/skills/`. Existing locked skills will be reported as drifted on the first sync after upgrade, and `thv skill sync --check` exits non-zero in CI until re-synced.

Sync skills only to a specific client to avoid unexpected drift into `.qoder/skills/` after the new `qoder` client was added.

```
$ thv skill sync --clients claude-code
thv skill sync --check --clients claude-code
```

— Names client, path, flags, and CI impact with example commands v0.45.0

06 **Validation for runtime_config.build_with on workloads API** 75

BREAKING

`runtime_config.build_with` on `npx://` or `go://` images submitted via `POST /api/v1beta/workloads` now returns `400 Bad Request` instead of silently discarding the field.

— Names exact endpoint, field, and status code change v0.45.0

07 **Sigstore verification for plugin artifacts** NEW 65

ToolHive now performs end-to-end Sigstore bundle verification for plugin artifacts; stored bundles and git commit payloads/signatures are size-capped at 1 MiB and rejected with `422` above that limit.

— Names mechanism and limit but no command to invoke it v0.45.0

08 **Package name validation for npx/uvx/go references** BREAKING 60

Package names in `npx://`, `uvx://`, and `go://` references are now constrained to `[A-Za-z0-9@/:._+=~[]-]` at build time; references with other characters fail the build instead of being interpolated into generated Dockerfiles, preventing shell metacharacter injection.

— Names the exact regex and affected schemes but no example call v0.45.0

09 **Virtual MCP timeout enforcement and health propagation** 60

BREAKING

Virtual MCP now honours `operational.timeouts` and propagates backend health changes to live sessions; a configured value below 30s will now actually cut backend calls that previously silently received the 30s default.

— Names config key and before/after default but no usage example v0.45.0

THINNER COVERAGE BELOW

10 **UTF-8 BOM policy bypass fix** IMPROVED 50

A UTF-8 BOM in `tools/list`, `prompts/list`, or `resources/list` responses no longer bypasses Cedar policy or tool-filter authorization.

— Names affected endpoints and policy engine, no action steps v0.45.0

11 **thv llm proxy auth error correction** IMPROVED 50

The `thv llm` local proxy now returns `401 token_required` instead of `502 server_error` when the stored credential has been rejected by the IdP.

— Names before/after status codes but no reproduction steps v0.45.0

12 **Non-2xx list response forwarding fix** IMPROVED 35

Non-2xx `list` responses from backends are now correctly forwarded with their original status code, instead of being delivered as HTTP 200 with an unfiltered body.

— Clear before/after but no named endpoint or config v0.45.0

13 **RFC 7523 flows in embedded auth server** NEW 20

Adds two new RFC 7523 flows to the embedded auth server.

— Single-line mention with no mechanism or usage detail v0.45.0

└──▷ **BREAKING ON UPGRADE**

! The `operator-crds` chart must be updated to 0.46.0 before or together with the operator chart — a stale CRD silently prunes `caBundleRef` from applied resources instead of rejecting it.

! `thv serve` now requires `Content-Type: application/json` on state-changing requests that carry a body over TCP; callers without it receive `415 Unsupported Media Type`. UNIX-socket paths (`--socket=<path>`) are unaffected.

! `thv serve` now validates `Origin` against a loopback-only allowlist for TCP listeners; non-loopback binds log a `WARN` and pass through.

! Package names for `npx://`, `uvx://`, and `go://` references are now constrained to `[A-Za-z0-9@/:._+=~[]-]`; references containing other characters fail at build time instead of being interpolated into the Dockerfile.

! `thv skill sync` without `--clients` now targets every skill-supporting client; existing locked skills are reported as drifted on first sync after upgrade, and `thv skill sync --check` exits non-zero in CI until re-synced.

! `runtime_config.build_with` on `npx://` or `go://` images via `POST /api/v1beta/workloads` now returns `400 Bad Request` instead of silently discarding the field.

! `thv skill push` now returns 400 if both `--key` and `--no-sign` are specified; exactly one of `--key`, `--identity-token`, or `--no-sign` is required.

! Virtual MCP now honours `operational.timeouts`; a configured value below 30 s will now actually cut backend calls that previously silently got the 30 s default.

! The `thv llm` local proxy returns `401 token_required` instead of `502 server_error` when the stored credential has been rejected by the IdP.

! Exported Go interfaces `plugins.MaterializationAdapter`, `state.Store` writers, `storage.UpstreamTokenStorage`, and six function signatures gained required methods or changed signatures; no effect on the CLI, operator, wire protocol, or persisted state.

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ

2 RELEASES · 2026-08-26 → 2026-08-28 NOTES ↗

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

> SkillSpector added transitive scanning of referenced skills to catch supply-chain risk, new BH1–BH3 hook/permission findings, SC9 hidden-artifact detection, and EA5 external-model-selection findings, alongside a machine-readable `risk_assessment.max_issue_severity` JSON field, LLM sampling/output-language controls, and safer sandboxed file traversal; it also moved `langgraph-cli[inmem]` into an optional extra, a breaking change for LangGraph Studio users.

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

01 **Transitive scanning of referenced skills** NEW 88

Adds `--transitive` flag to opt into transitive scanning of referenced skills, with `--transitive-depth`, `--transitive-allow-prefix`, and `--transitive-deny-prefix` controls for bounded traversal; partial provider execution and traversal truncation are surfaced in results so incomplete deep scans cannot silently appear clean.

Scan a skill and all skills it references transitively, limiting traversal depth and restricting allowed prefixes, to catch supply-chain risks in the full dependency graph.

```
$ skillspector scan ./my-skill/ --transitive --transitive-depth 2
--transitive-allow-prefix https://github.com/trusted-org/ --
transitive-deny-prefix https://github.com/untrusted-org/
```

— Exact flags with a runnable example and truncation caveat. v2.10.0

02 **langgraph-cli moved to optional extra** BREAKING 78

`langgraph-cli[inmem]` is moved from the base installation to the `langgraph-dev` optional extra; LangGraph Studio users who install only the base package must now install `skillspector[langgraph-dev]` explicitly.

— Exact package names with the required migration command. v2.10.0

03 **Machine-readable max issue severity in JSON report** NEW 76

Adds `risk_assessment.max_issue_severity` field to JSON output (value `NONE` when no active issue is reported) for machine-readable downstream policy gates.

Gate a CI pipeline on the highest severity found — fail the build if any HIGH or CRITICAL issue is present — by reading `risk_assessment.max_issue_severity` from the JSON report.

```
$ skillspector scan ./my-skill/ --format json --output
report.json && python -c "import json,sys;
r=json.load(open('report.json')); sys.exit(0 if
r['risk_assessment']['max_issue_severity'] in ('NONE','LOW') else
1)"
```

— Exact field name paired with a working CI gate example. v2.10.0

04 LLM provider configuration: sampling, language, listing NEW

75

Adds `SKILLSPECTOR_TEMPERATURE` (values `0` – `1`) and `SKILLSPECTOR_SEED` (integers) environment variables for optional LLM sampling control, forwarded to OpenAI-compatible and Azure OpenAI endpoints, with provider defaults preserved if unset; adds `SKILLSPECTOR_OUTPUT_LANGUAGE` to set the language of human-readable LLM-generated finding text; and expands `skillspector scan --help` to list all supported hosted, local, compatible, and CLI-backed LLM providers with their authentication paths.

Pin LLM output for reproducible scans when using an OpenAI-compatible endpoint — useful for CI pipelines where you want consistent results across runs.

```
$ SKILLSPECTOR_TEMPERATURE=0 SKILLSPECTOR_SEED=42 skillspector
scan ./my-skill/
```

Pass LLM sampling controls via a `.env` file for Docker-based scans to keep credentials and tuning out of shell history.

```ini
SKILLSPECTOR_PROVIDER=openai
OPENAI_API_KEY=sk-...
SKILLSPECTOR_TEMPERATURE=0
SKILLSPECTOR_SEED=42
```

Produce LLM finding text in French for a report shared with a francophone security team.

```
$ SKILLSPECTOR_OUTPUT_LANGUAGE=French skillspector scan ./my-skill/ --format markdown --output report.md
```

— Names every env var and flag with runnable examples. v2.11.0 v2.10.0

05 **BH1–BH3 hook and permission findings** BREAKING

75

Introduces BH1–BH3 findings by analyzing `hooks/hooks.json`, `.claude/settings.json`, and `.claude/settings.local.json` for bundled lifecycle hook execution (BH1), directly proven remote transfer of sensitive event or file content (BH2), and broad or ignored project permission surfaces (BH3). Existing scans may now surface these new findings; review them before relying on the existing baseline mechanism to suppress previously clean results.

— Exact files and finding codes named, but no scan command shown. v2.11.0

06 **SC9 hidden file and nested archive inspection** NEW

60

Adds bounded local inspection of hidden files and ZIP-compatible nested artifacts (DOCX, XLSX, PPTX) without extracting or executing members, raising HIGH SC9 findings for concealed executables.

— Names formats and finding code but no invocation shown. v2.10.0

THINNER COVERAGE BELOW

07 **Safe path traversal in restricted sandboxes** IMPROVED

53

Supports safe file traversal in restricted Linux sandboxes by opening intermediate path components with `O_PATH` where available, without requiring read access to every ancestor directory.

— Explains mechanism but has no user-facing command. v2.11.0

08 **EA5 external model/provider selection findings** NEW

48

Adds EA5 static findings for external model or provider selection, including silent coding-CLI account switches and top-level model pins.

— Describes scope but no example or command given. v2.10.0

09 **Structured skill summaries for AISOP/AISP bundles** NEW

46

Adds structured skill summaries for valid AISOP/AISP bundles across terminal, Markdown, JSON, and SARIF output formats.

— Names formats and bundle types but no command shown. v2.10.0

---

10 **Dynamic analyzer discovery and validation** NEW 30

Adds dynamic analyzer discovery and validates risk-score inputs against the registered analyzer set.

— Thin description with no named surface or example. v2.10.0

└──▷ **BREAKING ON UPGRADE**

! Existing scans may now report new BH1–BH3 findings for supported bundled hook and settings files (`hooks/hooks.json` , `.claude/settings.json` , `.claude/settings.local.json`); review those findings and use the existing baseline mechanism only after validating the declared execution or permission surface.

! `langgraph-cli[inmem]` is moved from the base installation to the `langgraph-dev` optional extra; LangGraph Studio users who install only the base package must now install `skillspector[langgraph-dev]` explicitly.

WAS THIS USEFUL? 👍 👎

NeMo Guardrails

Feedback ⓘ 1 RELEASE · 2026-08-26 NOTES ↗

NeMo Guardrails is an open-source toolkit for easily adding programmable guardrails to LLM-based conversational systems.

> NeMo Guardrails v0.24.0 introduces the RailOutcome contract for rail actions, typed rail manifests, and a major expansion of IORails to 59 of 67 built-in rail surfaces, alongside new server health endpoints, an F5 integration, and several breaking changes to action signatures, install extras, and message formats.

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

01 **RailOutcome contract for rail actions** BREAKING 88

Adds `RailOutcome`, an engine-neutral allow, block, or transform contract that rail actions return and that is enforced by both `LLMRails` and `IORails`. Custom actions previously using `@action(output_mapping=...)` must remove that argument and return an explicit `RailOutcome` instead.

— Names exact decorator argument and its replacement, clear migration path. v0.24.0

02 **Output-rail checking mode on /v1/checks** BREAKING 83

The `/v1/checks` endpoint gains an output-rail checking mode that selects a server-loaded configuration via `config_id` or the server default; it no longer accepts inline configuration.

— Names endpoint, parameter, and exact behavior change. v0.24.0

03 **IORails expanded to 59 of 67 rail surfaces** IMPROVED 83

Expands `IORails` to execute 59 of the library's 67 action-backed input and output surfaces, including community integrations and content-transforming rails, sharing implementations with `LLMRails`. Adds message checks to `IORails`, and non-streaming `generate()` and `generate_async()` calls now return `GenerationResponse` when generation options are supplied; message lists must be passed with `messages=` as a keyword argument rather than positionally.

— Names exact methods, return type, and new keyword requirement. v0.24.0

04 **hf-classifier install extra removed** BREAKING 70

The `hf-classifier` install extra is removed; install `transformers` and `torch` directly to use the local classifier backend.

— Exact extra name and replacement install steps given. v0.24.0

05 **snake_case action names for Colang 1 flows** BREAKING

62

Custom Colang 1 flows using the former space-separated Cleanlab, Fiddler, or GCP action names must switch to their `snake_case` names.

— Names affected integrations but not the exact old/new names. v0.24.0

06 **Server health-check endpoints** NEW

60

Adds `/v1/health` and `/healthz` server health-check endpoints.

Check server health before routing traffic to a NeMo Guardrails instance.

```
$ curl -sf http://localhost:8000/v1/health && echo 'server ready'
```

— Simple named endpoints with a runnable check example. v0.24.0

07 **Canonical outbound HTTP client** NEW

60

Adds a canonical outbound HTTP client shared by all built-in integrations, providing connection pooling, retries, TLS, error handling, lifecycle management, and privacy-safe tracing and metrics.

— Lists mechanisms but no config or invocation surface named. v0.24.0

THINNER COVERAGE BELOW

08 **Typed rail manifests for engine discovery** NEW

55

Adds typed rail manifests so built-in rails can declare their configuration, actions, execution surfaces, requirements, and privacy properties for automatic discovery by both `LLMRails` and `IORails`.

— Describes purpose but no manifest schema or fields given. v0.24.0

09 **Chat Completions message shape restrictions** BREAKING

55

Chat Completions requests must use supported role-specific OpenAI message shapes; internal event payloads, unexpected fields, and audio requests are now rejected.

— States new validation rules without listing exact schema. v0.24.0

10 **Multiple namespaced self-check rails** NEW

53

Adds support for multiple namespaced self-check rails, each with per-rail namespaced task prompts, enabling more than one self-check rail to run within the same configuration.

— Explains mechanism but no config keys or prompt names given. v0.24.0

11 **Model telemetry and instrumented decorator** NEW 35

Adds shared model telemetry and an instrumented model decorator for LLM call observability.

— Names decorator concept but gives no usage or config. v0.24.0

12 **F5 Guardrails integration** NEW 25

Adds F5 Guardrails integration to the built-in rail library.

— Names the integration only, no configuration detail given. v0.24.0

└──▷ **BREAKING ON UPGRADE**

! Custom actions using @action(output_mapping=...) must remove that argument and return an explicit `RailOutcome` instead.

! The `hf-classifier` install extra is removed; install `transformers` and `torch` directly to use the local classifier backend.

! `IORails` calls that supply generation options now return `GenerationResponse`; message lists must be passed with `messages=` as a keyword argument rather than positionally.

! Custom Colang 1 flows using the former space-separated Cleanlab, Fiddler, or GCP action names must switch to their `snake_case` names.

! `/v1/checks` no longer accepts inline configuration; select a server-loaded configuration via `config_id` or rely on the server default.

! Chat Completions requests must use supported role-specific OpenAI message shapes; internal event payloads, unexpected fields, and audio requests are now rejected.

WAS THIS USEFUL? 👍 👎

AI-Infra-Guard

Feedback ⓘ 1 RELEASE · 2026-08-26 NOTES ↗

A full-stack AI Red Teaming platform securing AI ecosystems via Agent Scan, Skills Scan, MCP scan, AI Infra scan and LLM jailbreak evaluation.

> AI-Infra-Guard v4.6.0 introduces an API security audit module with LLM API poisoning detection, overhauls its agent red-teaming mutation engine, and tightens vulnerability rule validation with fresh rule-set updates.

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

01 Mutation engine refactor for agent red-teaming BREAKING

70

The `aig-agent-redteam` v5.0.0 mutation engine refactor merges the `workflow-attack` mode into `mutation-attack`; any automation invoking `workflow-attack` will need to be updated to use `mutation-attack` instead.

– Names exact subcommands merged but lacks new behavior detail. v4.6.0

02 Vulnerability rule validation and rule-set updates IMPROVED

60

Adds strict YAML validation enforcing required `id` and `severity` fields in vulnerability rules, and updates AIG vulnerability rules with two new rule-set snapshots dated 2026-08-07 and 2026-08-14.

– Names required fields and exact snapshot dates but no usage steps. v4.6.0

THINNER COVERAGE BELOW

03 API security audit module with LLM API poisoning detection NEW

25

Adds a new API security audit module and detection for LLM API poisoning attacks.

– No mechanism or scope given beyond the module name. v4.6.0

04 DeepSeek Harness prompt injection research NEW

10

Includes DeepSeek Harness prompt injection assessment research.

– Bare mention with no detail on scope or findings. v4.6.0

└──▷ **BREAKING ON UPGRADE**

! The `workflow-attack` subcommand/mode in `aig-agent-redteam` is merged into `mutation-attack`; any automation invoking `workflow-attack` will need to be updated.

WAS THIS USEFUL?  

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Microsoft code-oss

Feedback ⓘ 1 RELEASE · 2026-08-26 BLOG ↗

Visual Studio Code

> VS Code introduces the Agent Host architecture: a standalone, persistent process for agent sessions accessible via a new open protocol, remote connections, a synchronized Agents window, and support for multiple agent harnesses.

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

01 Standalone Agent Host process NEW

75

A new `code agent host` (or `code-insiders agent host`) subcommand starts a standalone Agent Host process that keeps agent sessions alive independent of any editor window, making the process persistent and portable.

Start a standalone Agent Host so external AHP clients can connect to and monitor agent sessions running on your machine.

```
$ code agent host
```

— Names exact command; mechanism explained but scope kept brief launch-20260826-bf70a8c8

02 Agent Host Protocol and multi-language client libraries NEW

75

The Agent Host Protocol (AHP) is an open, URI-addressable protocol for coordinating multiple clients around the same long-running agent session, with channels for sessions, chats, terminals, and changesets. Client libraries ship in Rust (`ahp`, `ahp-types`, `ahp-ws`), TypeScript (`@microsoft/agent-host-protocol`), Kotlin (`com.microsoft.agenthostprotocol:agent-host-protocol`), Go (`github.com/microsoft/agent-host-protocol/clients/go`), and Swift (via Swift Package Manager: `microsoft/agent-host-protocol`).

— Names protocol channels and every client library package verbatim

launch-20260826-bf70a8c8

03 Remote agent sessions over SSH, tunnels, and web NEW

75

The Agent Host can run on a remote machine while clients connect from the desktop or browser via SSH, dev tunnels, or the web at `insiders.vscode.dev/agents`, letting practitioners review and manage a session running elsewhere.

Connect to a remote Agent Host from the browser to review and manage a session running on another machine.

> 📍 1. On the remote machine, start VS Code with your workspace and begin an agent session. 2. In the Agents window, connect to the remote Agent Host (configured via SSH or dev tunnel). 3. Navigate to `insiders.vscode.dev` /agents in a browser, select the same host, and continue the session.

— Names exact connection methods and web URL with worked steps `launch-20260826-bf70a8c8`

---

04 **Multi-harness agent sessions (Copilot, Claude)** NEW

60

Multiple agent harnesses can run under a common session model: Copilot via the GitHub Copilot SDK and Claude via the Anthropic Claude Agent SDK, with each harness retaining its own agent loop, tools, and provider-specific capabilities.

.

— Names both SDKs and harness picker but no exact command `launch-20260826-bf70a8c8`

THINNER COVERAGE BELOW

05 **Persistent agent sessions across window closes** NEW

55

Agent sessions persist after closing a folder or editor window, with state and progress resuming when the folder is reopened, so work is not lost when a window is closed rather than quit.

Experience a persistent session that survives closing its folder: start a session, close the folder, then reopen it or switch to the Agents window to resume.

> 📍 1. Open a folder in an editor window, select Copilot from the harness picker, and start an agent session. 2. While the agent is working, close the folder (do not quit VS Code). 3. Reopen the folder — session state and progress are preserved. 4. Open the Agents window to monitor the same live session from both surfaces simultaneously.

— Behavior explained with a worked example but no named config `launch-20260826-bf70a8c8`

---

06 **Synchronized Agents window** NEW

45

A new Agents window acts as a synchronized client surface, letting practitioners monitor and interact with a live agent session from both the editor window and the Agents window simultaneously without creating a copy.

Experience a persistent session that survives closing its folder: start a session, close the folder, then reopen it or switch to the Agents window to resume.

> 📍 1. Open a folder in an editor window, select Copilot from the harness picker, and start an agent session. 2. While the agent is working, close the folder (do not quit VS Code). 3. Reopen the folder — session state and progress are preserved. 4. Open the Agents window to monitor the same live session from both surfaces simultaneously.

— UI surface named but reached only via prose, no exact path `launch-20260826-bf70a8c8`

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Cloudflare is a web performance and security platform providing DDoS protection, CDN, DNS, and developer tools.

> Cloudflare's biggest change this window is a major API refresh (v4.0.0) adding 230 endpoints and removing 142 — new Agent Memory, Cloudforce One rules, Data Security posture, Email Security content-policy, Images, Logpush transformer and expanded billing APIs, a Secrets Store path change, and a large Workers AI model catalog pruning — alongside symmetric JWT key support, Durable Objects concurrency and alarm-control improvements, new Workers AI models including a multimodal GLM-5.3 Flash, Log Explorer dataset deletion, scannable Access token secrets, and a hard 2026-09-14 deadline to migrate off the Azure-based Microsoft Sentinel connector.

└──▷ WHAT SHIPPED · 19 FEATURES most completely described first ? what's the number?

01 **Log Explorer dataset deletion** NEW 86

Log Explorer datasets can now be permanently deleted: set the new `deletion_protection` field to `false` via the Update an account or zone dataset API method (`PUT /accounts/{account_id}/logs/explorer/datasets/{dataset_id}` or `PUT /zones/{zone_id}/logs/explorer/datasets/{dataset_id}`), then call `DELETE /accounts/{account_id}/logs/explorer/datasets/{dataset_id}` or `DELETE /zones/{zone_id}/logs/explorer/datasets/{dataset_id}` to remove it from the dashboard or API.

— Exact field and both delete endpoints given 4.0.0

02 **Durable Object alarm abort control** NEW 85

ctx.abort() on Durable Objects now accepts a `{ retryAlarm: false }` option so an interrupted alarm stops retrying instead of restarting after reset; requires Wrangler 4.126.0 or later for local dev.

Stop a Durable Object alarm from retrying when the alarm's own cleanup work is complete, so the alarm handler does not re-run unnecessarily.

```typescript
import { DurableObject } from 'cloudflare:workers';

export class CleanupTask extends DurableObject {
  async alarm(): Promise<void> {
    await this.ctx.storage.deleteAll();
    this.ctx.abort('Cleanup complete', { retryAlarm: false });
  }
}
```

— Runnable example with exact option and version requirement product docs

03 **GLM-5.3 Flash multimodal model on Workers AI** NEW 83

`@cf/zai-org/glm-5.3-flash` (GLM-5.3 Flash) is now available on Workers AI as the first natively multimodal GLM-5 model, with 320B total parameters (18B active per token on MoE) and a 1,048,576-token context window; accessible via env.AI.run(), the REST API, or the OpenAI-compatible endpoint, and requires prepaid AI Gateway credits.

Run the new multimodal GLM-5.3 Flash model on Workers AI for a text generation task with its large 1M-token context window.

```javascript
const response = await env.AI.run('@cf/zai-org/glm-5.3-flash', {
  messages: [{ role: 'user', content: 'Summarize this incident
report: ...' }]
});
```

— Mechanism, parameters and runnable example given product docs

04 **Agent Memory API for AI agents** NEW 80

New Agent Memory API for storing and retrieving AI agent context: create and manage namespaces (`POST /accounts/{account_id}/agent-memory/namespaces` , `GET /accounts/{account_id}/agent-memory/namespaces` , `GET /accounts/{account_id}/agent-memory/namespaces/{namespace_name}` , `DELETE /accounts/{account_id}/agent-memory/namespaces/{namespace_name}`), manage per-profile memories and sessions (`GET .../profiles/{profile_name}/memories` , `GET .../memories/{memory_id}` , `DELETE .../profiles/{profile_name}` , `DELETE .../memories/{memory_id}` , `DELETE .../sessions/{session_id}`), and ingest, recall, remember or summarize content (`POST .../ingest` , `POST .../recall` , `POST .../remember` , `POST .../summary`).

— Endpoints named but no request/response schema shown 4.0.0

05 **Cloudforce One structured rules and approvals** NEW

80

New structured detection-rule workflow in Cloudforce One: create, validate, test and update rules (`GET /accounts/{account_id}/cloudforce-one/rules/structured` , `GET .../structured/schema` , `GET .../structured/{id}` , `POST .../structured` , `POST .../structured/validate` , `POST .../structured/{id}/test` , `PUT .../structured/{id}`), manage approvals (`GET .../rules/approvals` , `GET .../approvals/{id}` , `POST .../approvals/{id}/cancel` , `POST .../approvals/{id}/resubmit` , `PUT .../structured/approvals/{id}`), and manage rule exemptions (`GET` , `POST` , `PUT` , `DELETE` `/accounts/{account_id}/cloudforce-one/rules/exemptions`).

— Full endpoint set named, no schema detail `4.0.0`

---

06 **Billing and payment methods API expansion** NEW

77

Expanded account billing API: full CRUD for payment methods (`GET/POST /accounts/{account_id}/payment-methods` , `GET/PUT/DELETE .../payment-methods/{payment_method_id}` , `POST .../payment-methods/{payment_method_id}/set-as-default`) and billing profile (`GET/POST/PUT/DELETE /accounts/{account_id}/billing/profile`), new read endpoints for credits, history, bad debt and unpaid invoices (`GET .../billing/credits` , `.../billing/history` , `.../billing/bad-debt` , `.../billing/unpaid-invoice`), new actions to pay bad debt or an invoice (`POST .../pay-bad-debt` , `POST .../pay-invoice`), `PATCH .../invoices` , invoice PDF retrieval (`GET .../receipts/{receipt_id}/pdf`), and address validation (`POST /billing/address-validation`).

— Many endpoints named, no request payload detail `4.0.0`

---

07 **Data Security posture policy API** NEW

75

New Data Security posture policy management: list and create policies (`GET/POST /accounts/{account_id}/data-security/posture/policies`), get/update/delete a policy (`GET/PUT/DELETE .../policies/{policy_id}`), view policy logs (`GET .../policies/logs`), and ignore or unignore findings (`POST .../findings/ignore` , `POST .../findings/unignore`).

— Endpoints named, no policy schema detail `4.0.0`

---

08 **Email Security content policy management** NEW

73

New content-policy management for Email Security: list, create, batch-create, update and delete content policies (`GET/POST /accounts/{account_id}/email-security/settings/content_policies` , `GET/PATCH/DELETE .../content_policies/{policy_id}` , `POST .../content_policies/batch`), plus new endpoints to update allow, block and trusted-domain rules (`PUT`

`.../allow_policies/{policy_id}`, `PUT .../block_senders/{pattern_id}`, `PUT .../trusted_domains/{trusted_domain_id}`).

— Endpoints named, no field-level schema given `4.0.0`

09 **Images v1/v2 API** NEW

73

New Images v1 and v2 API surface: full CRUD on v1 images (`GET/POST /accounts/{accountId}/v1/images`, `GET/PUT/PATCH/DELETE .../v1/images/{imageId}`, `GET .../v1/images/{imageId}/blob`), a v2 listing endpoint (`GET /accounts/{accountId}/v2/images`) and direct upload (`POST /accounts/{accountId}/v2/images/direct_upload`), plus zone image flows (`GET/PUT /accounts/{accountId}/zones/{zoneId}/v1/images/flows`).

— Endpoints named, no upload payload spec given `4.0.0`

10 **Secrets Store endpoint path restructuring** BREAKING

73

Secrets Store endpoints no longer take `account_id`, `store_id` or `secret_id` as path parameters on `DELETE/GET/POST/PATCH /accounts/{account_id}/secrets_store/stores`, `.../stores/{store_id}`, `.../stores/{store_id}/secrets`, `.../stores/{store_id}/secrets/{secret_id}`, `.../secrets/{secret_id}/duplicate` and `.../secrets_store/quota`, changing how these routes must be called.

— Named endpoints changed, no migration guidance given `4.0.0`

11 **New Workers AI text generation models in AI Search** NEW

65

AI Search now supports six additional text generation models on Workers AI: `@cf/deepseek-ai/deepseek-v4-flash-0731` and `@cf/deepseek-ai/deepseek-v4-pro-0813` (1,048,576-token context), `@cf/openai/gpt-oss-120b` and `@cf/openai/gpt-oss-20b` (128,000-token context), `@cf/qwen/qwen3.8-27b` (262,144-token context), and `@cf/moonshotai/kimi-k2.7-code` (262,144-token context).

— Model names and context windows listed, no usage example `product docs`

12 **Scannable Access service token secret format** IMPROVED

63

Access service token Client Secrets created on or after August 26, 2026 use the scannable format `cfast_[40 alphanumeric characters][8-character checksum]`, making them detectable by secret-scanning tools; existing `CF-Access-Client-Id` / `CF-Access-Client-Secret` headers are unchanged.

— Exact format and cutover date given, no action required product docs

13 **Symmetric HMAC keys for API Shield JWT validation** NEW

62

API Shield JWT validation now supports symmetric HMAC keys (`HS256`, `HS384`, `HS512`), configurable via the Cloudflare dashboard or API; credentials are never stored in plaintext and are never returned in API responses.

— Algorithms named, no exact config endpoint given

API Shield - Symmetric key support for …

14 **Workers AI model catalog pruning** BREAKING

60

Cloudflare removed run endpoints for dozens of Workers AI models, including `@cf/meta/llama-3.1-70b-instruct-fp8-fast`, `@cf/meta/llama-4-scout-17b-16e-instruct`, `@cf/mistral/mistral-7b-instruct-v0.1`, `@cf/qwen/qwen3-30b-a3b-fp8`, `@cf/openai/whisper`, `@cf/openai/gpt-oss-120b`, `@cf/black-forest-labs/flux-1-schnell`, `@cf/stabilityai/stable-diffusion-xl-base-1.0` and `@cf/baai/bge-large-en-v1.5`, among many others, breaking any code calling `POST /accounts/{account_id}/ai/run/{model}` for those model paths.

— Sample of removed models given, full list omitted 4.0.0

15 **Microsoft Sentinel connector migration deadline** DEPRECATED

60

The Azure Functions-based Microsoft Sentinel Logpush connector is deprecated in favor of the Cloudflare for Microsoft Sentinel Codeless Connector Framework (CCF) connector; Enterprise customers must migrate by 2026-09-14, when the underlying Microsoft Azure Monitor HTTP Data Collector API is retired and Cloudflare stops maintaining the old connector.

— Deadline and replacement named, no migration steps given product docs

THINNER COVERAGE BELOW

16 **Legacy suppression and analytics endpoints deprecated**

DEPRECATED

55

Several legacy endpoints are deprecated but still function: email routing/sending suppression (`GET/POST/DELETE /accounts/{account_id}/email/routing/suppression`, `.../routing/suppression/{suppression_id}`, `/accounts/{account_id}/email/sending/suppression`, `.../sending/suppression/{suppression_id}`, and their zone-level equivalents under `/zones/{zone_id}/email/...`), plus DNS and rate-limit analytics (`GET /accounts/{account_id}/dns_firewall/{dns_firewall_id}/dns_analytics/report`, `.../dns_analytics/report/bytime`, `GET /zones/{zone_id}/dns_analytics/report`, `.../report/bytime`, and `GET /zones/{zone_id}/rate_limit_analytics`).

— Endpoints named, no replacement or sunset date given `4.0.0`

---

17 **Radar Researcher connected insights and AI explanations**

IMPROVED 50

Radar Researcher (beta) now surfaces connected insights linking to Radar pages, reports and Cloudflare Blog posts, adds an Explain with AI button on URL Scanner reports for AI-driven findings explanations and follow-up Q&A, and lets users open shared sessions in fullscreen.

— UI features described, no exact navigation path given `product docs`

---

18 **Durable Objects Dynamic Workers concurrency raised to ten**

IMPROVED 45

Each Durable Object can now use up to ten concurrent Dynamic Workers with in-flight requests, up from four, raising the parallelism ceiling for stateful workloads.

— Numeric before/after given, no config knob `product docs`

---

19 **AI Search metadata envelope increase** IMPROVED 45

AI Search now supports larger custom metadata values within a shared 10 KiB metadata envelope per vector, with the first 64 UTF-8 bytes of each indexed string remaining filterable.

— Limit stated, no config key or example given `product docs`

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! Removed endpoint DELETE `/accounts` /{account_id} `/cloudforce-one` /exemptions

! Removed endpoint DELETE `/accounts` /{account_id} `/containers` /applications/{application_id} `/instances` /{instance_id}

! Removed endpoint DELETE `/system` /accounts/{account_tag} `/stores` /{store_id}

! Removed endpoint DELETE `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets`

! Removed endpoint DELETE `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets` /{secret_id}

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-1

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-1-internal

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2-en

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2-en-ws

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2-es

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/flux

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/nova-3

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/nova-3-internal

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/deepgram/nova-3-ws

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/nvidia/nemotron-speech-streaming-en-0.6b

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/pipecat-ai/smart-turn-v2

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/pipecat-ai/smart-turn-v3

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/sven/test-pipe-http

! Removed endpoint GET `/accounts` /{account_id} `/ai` /run/@cf/test/hello-world-cog

! Removed endpoint GET `/accounts` /{account_id} `/cloudforce-one` /exemptions

! Removed endpoint GET `/accounts` /{account_id} `/containers` /instances/{instance_id} `/ssh`

! Removed endpoint GET `/accounts` /{account_id} `/email` /sending/feedback

! Removed endpoint GET `/accounts` /{account_id} `/paygo-usage`

! Removed endpoint GET `/accounts` /{account_id} `/paygo-usage-info`

! Removed endpoint GET `/system` /accounts/{account_tag} `/stores`

! Removed endpoint GET `/system` /accounts/{account_tag} `/stores` /{store_id}

! Removed endpoint GET `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets`

! Removed endpoint GET `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets` /{secret_id}

! Removed endpoint PATCH `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets` /{secret_id}

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/ai4bharat/indictrans2-en-indic-1B

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/ai4bharat/nonomni-indictrans2-en-indic-1b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/aisingapore/gemma-sea-lion-v4-27b-it

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/bge-base-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/bge-large-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/bge-m3

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/bge-reranker-base

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/bge-small-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/nonomni-bge-base-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/nonomni-bge-large-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/nonomni-bge-m3

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/baai/nonomni-bge-small-en-v1.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/black-forest-labs/flux-1-schnell

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/black-forest-labs/flux-2-dev

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/black-forest-labs/flux-2-klein-4b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/black-forest-labs/flux-2-klein-9b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/bytedance/stable-diffusion-xl-lightning

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-1

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2-en

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepgram/aura-2-es

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepgram/flux

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepgram/nova-3

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepseek-ai/deepseek-math-7b-instruct

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/deepseek-ai/deepseek-r1-distill-qwen-32b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/defog/sqlcoder-7b-2

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/facebook/bart-large-cnn

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/facebook/nonomni-bart-large-cnn

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/facebook/nonomni-detr-resnet-50

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/fblgit/una-cybertron-7b-v2-bf16

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/embeddinggemma-300m

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/gemma-2b-it-lora

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/gemma-3-12b-it

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/gemma-4-26b-a4b-it

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/gemma-7b-it-lora

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/google/nonomni-embeddinggemma-300m

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/huggingface/distilbert-sst-2-int8

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/huggingface/nonomni-distilbert-sst-2-int8

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/ibm-granite/granite-4.0-h-micro

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/leonardo/lucid-origin

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/leonardo/phoenix-1.0

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/lykon/dreamshaper-8-lcm

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta-llama/llama-2-7b-chat-hf-lora

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-2-7b-chat-fp16

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-2-7b-chat-int8

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3-8b-instruct

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3-8b-instruct-awq

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.1-70b-instruct-fp8-fast

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.1-8b-instruct-awq

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.1-8b-instruct-fp8

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.1-8b-instruct-fp8-fast

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.2-11b-vision-instruct

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.2-1b-instruct

! Removed endpoint POST `/accounts`/{account_id}`/ai`/run/@cf/meta/llama-3.2-3b-instruct

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/meta/llama-3.3-70b-instruct-fp8-fast

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/meta/llama-4-scout-17b-16e-instruct

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/meta/llama-guard-3-8b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/meta/m2m100-1.2b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/microsoft/nonomni-resnet-50

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/microsoft/phi-2

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/microsoft/resnet-50

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/mistral/mistral-7b-instruct-v0.1

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/mistral/mistral-7b-instruct-v0.2-lora

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/mistralai/mistral-small-3.1-24b-instruct

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/moonshotai/kimi-k2.5

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/moonshotai/kimi-k2.6

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/moonshotai/kimi-k2.7-code

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/myshell-ai/melotts

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/nvidia/nemotron-3-120b-a12b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/nvidia/nemotron-speech-streaming-en-0.6b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openai/gpt-oss-120b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openai/gpt-oss-20b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openai/whisper

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openai/whisper-large-v3-turbo

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openai/whisper-tiny-en

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/openchat/openchat-3.5-0106

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/pfnet/plamo-embedding-1b

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen1.5-0.5b-chat

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen1.5-1.8b-chat

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen1.5-14b-chat-awq

! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen1.5-7b-chat-awq

- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen2.5-coder-32b-instruct
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen3-30b-a3b-fp8
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwen3-embedding-0.6b
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/qwen/qwq-32b
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/runwayml/stable-diffusion-v1-5-img2img
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/runwayml/stable-diffusion-v1-5-inpainting
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/stabilityai/stable-diffusion-xl-base-1.0
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/thebloke/discolm-german-7b-v1-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/tiiuae/falcon-7b-instruct
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/tinyllama/tinyllama-1.1b-chat-v1.0
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/zai-org/glm-4.7-flash
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@cf/zai-org/glm-5.2
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/google/gemma-7b-it
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/mistral/mistral-7b-instruct-v0.2
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/nexusflow/starling-lm-7b-beta
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/nousresearch/hermes-2-pro-mistral-7b
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/deepseek-coder-6.7b-base-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/deepseek-coder-6.7b-instruct-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/llama-2-13b-chat-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/mistral-7b-instruct-v0.1-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/neural-chat-7b-v3-1-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/openhermes-2.5-mistral-7b-awq
- ! Removed endpoint POST `/accounts` /{account_id} `/ai` /run/@hf/thebloke/zephyr-7b-beta-awq

! Removed endpoint POST `/accounts` /{account_id} `/cloudforce-one` /exemptions

! Removed endpoint POST `/accounts` /{account_id} `/containers` /applications/{application_id} `/instances`

! Removed endpoint POST `/accounts` /{account_id} `/containers` /applications/{application_id} `/instances` /{instance_id} `/exec`

! Removed endpoint POST `/accounts` /{account_id} `/containers` /applications/{application_id} `/instances` /{instance_id} `/fetch`

! Removed endpoint POST `/system` /accounts/{account_tag} `/stores`

! Removed endpoint POST `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets`

! Removed endpoint POST `/system` /accounts/{account_tag} `/stores` /{store_id} `/secrets` /{secret_id} `/duplicate`

! Removed endpoint PUT `/accounts` /{account_id} `/intel` /attack-surface-report/{issue_id} `/dismiss`

! DELETE `/accounts` /{account_id} `/pipelines` /v1/sinks/{sink_id}: param force (query) removed

! DELETE `/accounts` /{account_id} `/pipelines` /v1/streams/{stream_id}: param force (query) removed

! DELETE `/accounts` /{account_id} `/secrets` _store/stores/{store_id}: param account_id (path) removed; param store_id (path) removed

! DELETE `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` : param account_id (path) removed; param store_id (path) removed; request body changed

! DELETE `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` /{secret_id}: param account_id (path) removed; param secret_id (path) removed; param store_id (path) removed

! GET `/accounts` /{account_id} `/cloudforce-one` /events: param cache (query) added; param source (query) removed

! GET `/accounts` /{account_id} `/cloudforce-one` /events/indicators: param cache (query) added; param cursor (query) added; param source (query) removed; response schema changed

! GET `/accounts` /{account_id} `/containers` /applications: param label (query) removed

! GET `/accounts` /{account_id} `/email` /sending/suppressions: param search (query) added; param direction (query) removed; param order (query) removed; response schema changed

! GET `/accounts` /{account_id} `/secrets` _store/quota: param account_id (path) removed

! GET `/accounts` /{account_id} `/secrets` _store/stores: param account_id (path) removed

! GET `/accounts` /{account_id} `/secrets` _store/stores/{store_id}: param account_id (path) removed; param store_id (path) removed

! GET `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` : param account_id (path) removed; param store_id (path) removed

! GET `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` /{secret_id}: param account_id (path) removed; param secret_id (path) removed; param store_id (path)

removed

! PATCH `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` /{secret_id}: param account_id (path) removed; param secret_id (path) removed; param store_id (path) removed; request body changed

! POST `/accounts` /{account_id} `/secrets` _store/stores: param account_id (path) removed

! POST `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` : param account_id (path) removed; param store_id (path) removed; request body changed

! POST `/accounts` /{account_id} `/secrets` _store/stores/{store_id} `/secrets` /{secret_id} `/duplicate` : param account_id (path) removed; param secret_id (path) removed; param store_id (path) removed; request body changed

! Deprecated endpoint DELETE `/accounts` /{account_id} `/email` /routing/suppression/{suppression_id} (still works)

! Deprecated endpoint DELETE `/accounts` /{account_id} `/email` /sending/suppression/{suppression_id} (still works)

! Deprecated endpoint DELETE `/zones` /{zone_id} `/email` /routing/suppression/{suppression_id} (still works)

! Deprecated endpoint DELETE `/zones` /{zone_id} `/email` /sending/suppression/{suppression_id} (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/dns` _firewall/{dns_firewall_id} `/dns` _analytics/report (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/dns` _firewall/{dns_firewall_id} `/dns` _analytics/report/bytime (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/email` /routing/suppression (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/email` /routing/suppression/{suppression_id} (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/email` /sending/suppression (still works)

! Deprecated endpoint GET `/accounts` /{account_id} `/email` /sending/suppression/{suppression_id} (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/dns` _analytics/report (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/dns` _analytics/report/bytime (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/email` /routing/suppression (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/email` /routing/suppression/{suppression_id} (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/email` /sending/suppression (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/email` /sending/suppression/{suppression_id} (still works)

! Deprecated endpoint GET `/zones` /{zone_id} `/rate` _limit_analytics (still works)

- ! Deprecated endpoint POST `/accounts`/{account_id}`/email`/routing/suppression (still works)
- ! Deprecated endpoint POST `/accounts`/{account_id}`/email`/sending/suppression (still works)
- ! Deprecated endpoint POST `/zones`/{zone_id}`/api`_gateway/expression-template/fallthrough (still works)
- ! Deprecated endpoint POST `/zones`/{zone_id}`/email`/routing/suppression (still works)
- ! Deprecated endpoint POST `/zones`/{zone_id}`/email`/sending/suppression (still works)
- ! Enterprise customers using the Azure Functions-based Microsoft Sentinel Logpush connector must migrate to the Cloudflare for Microsoft Sentinel CCF connector by 2026-09-14, when the underlying Microsoft Azure Monitor HTTP Data Collector API is retired and Cloudflare will stop maintaining the old connector.
- ! The Azure Functions-based Microsoft Sentinel Logpush Connector is deprecated; Cloudflare Enterprise customers must migrate to the Cloudflare for Microsoft Sentinel CCF connector by 2026-09-14.
- ! Cloudflare Enterprise customers using the Azure Functions-based Microsoft Sentinel Logpush connector must migrate to the Cloudflare for Microsoft Sentinel CCF connector by 2026-09-14, after which the old connector will be deprecated.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 28, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe** Feedback on this issue