THE CYBER TOOLCHAIN

NO. 047

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 16, 2026 · EVERY DAY

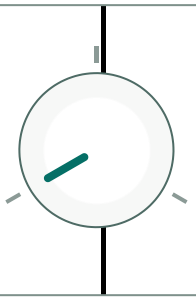
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 14 tools

APPLICATION SECURITY

vim ? help

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ 1 RELEASE · 2026-08-16

CHANGELOG ↗

GitGuardian SaaS is a cloud-based security platform that detects and prevents secret leaks in Git repositories and source code.

> GitGuardian SaaS added new API endpoints for secret location listing and live scan progress, plus a custom secret-validity API, honeytoken protection for developer endpoints, and programmatic PAT/SAT token management.

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

01 Secret location listing by incident and source NEW

95

New `GET /v1/incidents/secrets/{incident_id}/locations` endpoint lists every file on a source's default branch still holding a secret for a given incident, with exact line positions, filterable by `source`, `number_of_mentions_min`, `number_of_mentions_max`, `code_fix_request_status` (`processing`, `created`, `error`), and `search` (filepath substring), sortable by `last_updated_at`, `number_of_mentions`, or `source`. A companion `GET /v1/sources/{source_id}/locations` endpoint lists all files in a single source with any secret on the default branch, carrying the same per-match line positions, filters, and sort options.

Find every file still exposing the secret for a given incident, scoped to actively monitored sources with more than one mention, to prioritize remediation.

```
$ curl -s -H 'Authorization: Token <api_key>'
'https://api.gitguardian.com/v1/incidents/secrets/42/locations?
monitored=true&number_of_mentions_min=2&ordering=-
number_of_mentions'
```

List all secret locations in a specific source that match a filepath pattern, to hand off an exact file list to the owning team.

```
$ curl -s -H 'Authorization: Token <api_key>'
'https://api.gitguardian.com/v1/sources/5523/locations?
search=settings.py&code_fix_request_status=created'
```

— Full endpoint paths, filters, and runnable curl examples given. api-docs-983de2edf972

02 **Honeytoken protection for developer endpoints** NEW

93

The `ggshield honeytoken plant` command (via the existing agent binary) plants a decoy AWS credential on developer endpoints; any use of the decoy triggers an alert naming the compromised machine, its IP address, user agent, and timestamp, with no new agent required. A new `Honeytoken protection` card in the Endpoints dashboard reports fleet coverage percentage and lists unprotected endpoints.

Plant a honeytoken on a developer machine to detect credential-harvesting malware the moment it reads the credential file.

```
$ ggshield honeytoken plant
```

— Named command, alert fields, and dashboard card, with runnable example.

changelog-20260816-37c07038

03 **Live scan progress polling endpoint** NEW

90

New beta `GET /v1/scans/{scan_id}/live-progress` endpoint returns the live status (`queued` , `in_progress` , `completed`) and real-time counters (`elements_scanned` , `elements_total` , `elements_skipped` , `elements_failed` , `data_downloaded_bytes` , `data_scanned_bytes`) for a running historical scan, including an `audit_url` field once the scan completes.

Poll a running historical scan for live progress counters and stop when `completed` is returned, then fetch the audit summary from `audit_url` .

```
$ curl -s -H 'Authorization: Token <api_key>'
'https://api.gitguardian.com/v1/scans/<scan_id>/live-progress'
```

— Named endpoint and fields with a runnable poll example. api-docs-983de2edf972

04 **Custom secret validity submission via API** NEW

70

A new API endpoint lets teams push their own validity status (`valid` or `invalid`) for secrets on unreachable internal systems, authenticated with a personal access token scoped to `secrets:write` ; the status propagates to all incidents for that secret, appears as `Custom validity` with a datestamp, and is recorded in the incident activity log and audit log.

— Scope and effects named but exact endpoint path missing. changelog-20260816-37c07038

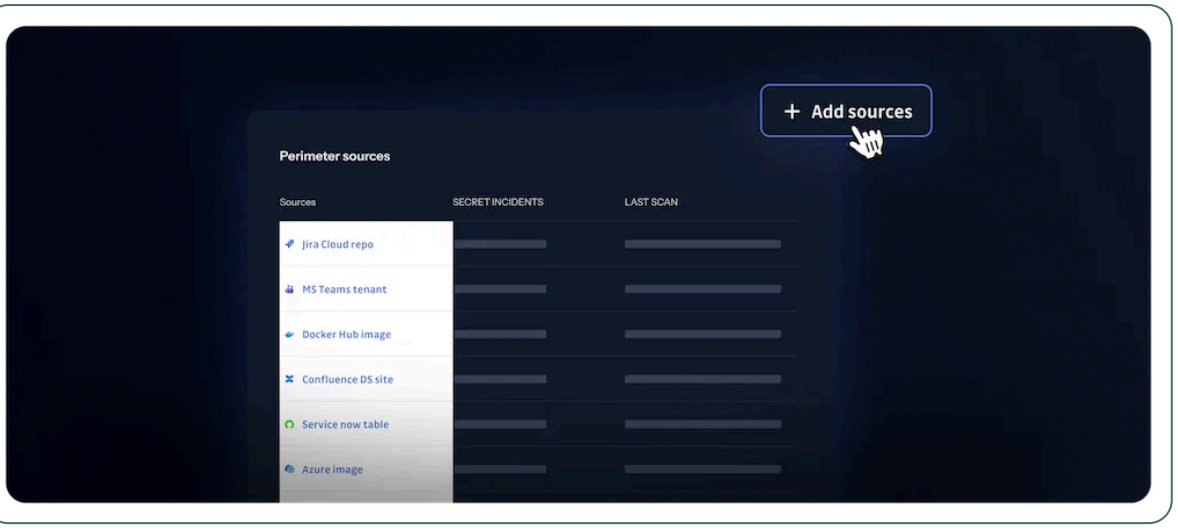
05 **Programmatic PAT/SAT creation and audit logging** NEW

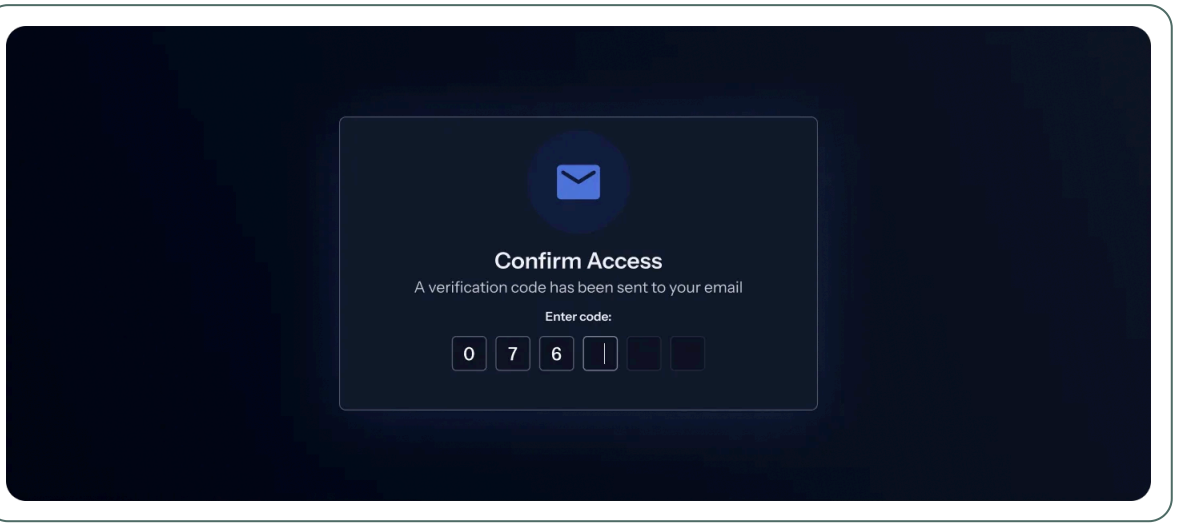
60

The API now supports creating Personal Access Tokens (PATs) and Service Account Tokens (SATs) programmatically for automated token rotation workflows; tokens created via the API must have a finite TTL. Scans triggered via a Personal Access Token now also emit an audit log event.

— Mechanism and TTL rule named but no endpoint path given. `changelog-20260816-37c07038`

└──▷ ALSO FROM THESE RELEASES





WAS THIS USEFUL?

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-08-14 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds advanced configuration options in this release.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-14
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-14
```

Adds advanced configuration options to extend sandbox behavior (see PR #3160 for details).

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ IDS / IPS

Fail2ban

Feedback ⓘ 1 RELEASE · 2026-08-15 NOTES ↗ CODE ↗

Daemon to ban hosts that cause multiple authentication errors

> Fail2ban 1.1.1 adds file-based IP allowlists with dynamic reload, a new OpenVPN jail and CSF firewall action, and extends the exim and sshd filters with new detection modes.

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

01 **File-based IP allowlists via ignoreip** NEW 95

Extends `ignoreip` with `file:...` (or `file://path/file-with-ip-set`) syntax to load ignored IPs, subnets, DNS/FQDN, or raw strings from an external file. The file is read lazily and reloaded dynamically when it changes, with comments after `#` or `;` ignored; implemented via the new `FileIPAddrSet` class in `ipdns.py`.

Maintain a shared allowlist file of trusted IPs and subnets that Fail2ban reloads automatically — useful for dynamic allowlists shared across jails.

ini

```
# In jail.local or jail.conf
[DEFAULT]
ignoreip = 127.0.0.1/8 ::1 file:///etc/fail2ban/trusted-ips.txt
```

— Full mechanism, exact syntax and reload behaviour with runnable config 1.1.1

02 **Extended filter modes for exim and sshd** IMPROVED 80

Adds `filter.d/exim.conf` mode `more` (several rules split out from mode `normal`) and extends mode `aggressive` to catch ACL-dropped failures such as `ACL: Country is banned`. Also extends `filter.d/sshd.conf` modes `ddos` and `aggressive` with regex for timeout-before-authentication connections.

Enable aggressive exim filtering to also block connections dropped by ACL rules (e.g. country bans) in addition to normal authentication failures.

ini

```
# In jail.local
[exim]
enabled = true
filter  = exim[mode=aggressive]
```

— Names exact filter files and modes with a config example 1.1.1

THINNER COVERAGE BELOW

03 **New OpenVPN jail and CSF firewall action** NEW 45

Adds a new OpenVPN jail (`filter.d/openvpn.conf`) and a new CSF (ConfigServer Firewall) action, comparable to the existing APF action.

– Names new filter and action but no usage detail 1.1.1

04 **systemd-journal support for vsftpd and lighttpd-auth filters** NEW 45

Adds systemd-journal regex support to `filter.d/vsftpd.conf` and `filter.d/lighttpd-auth.conf` .

– Names both filter files, no further mechanism given 1.1.1

05 **Observer module tracks manual ban attempts** IMPROVED 35

Makes manually added ban attempts inform the observer module, taking the known ban-count of bad IPs into account.

– Explains behaviour but no config surface or example 1.1.1

06 **Default mysql_log path updated** IMPROVED 30

Updates `paths-common.conf` to change the default `mysql_log` path, which is the default `logpath` for the `mysqld-auth` jail.

– Names config file and jail but not the new path value 1.1.1

07 **OpenRC init script starts after nftables** IMPROVED 30

Updates the OpenRC init script (`files/fail2ban-openrc.init.in`) to start after `nftables` in addition to `iptables` .

– Names exact file and firewalls affected, minimal detail 1.1.1

08 **Debian package enables systemd fail2ban.socket** NEW 30

Debian package now installs and enables a systemd-managed `fail2ban.socket` (gh-4220).

– Names the socket unit and issue reference, no usage detail 1.1.1

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Rspamd

Feedback ⓘ 1 RELEASE · 2026-08-15 NOTES ↗

Rapid spam filtering system.

> Rspamd 4.1.5 introduces fuzzy sender-fact sharing over encrypted fuzzy rules, font-aware PDF text extraction, and a restructured Elastic fuzzy-results schema, alongside a wave of breaking changes covering file/SHM input gating, the GPT condition callback, fuzzy storage ACLs and encryption enforcement, and new controller/proxy connection limits.

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

01 TCP gating for file and shared-memory message inputs

BREAKING

96

`allow_file_and_shm_inputs` is added explicitly to the normal, controller, and proxy worker configs (`local.d/worker-normal.inc` , `local.d/worker-controller.inc` , `local.d/worker-proxy.inc`) to gate TCP access to `File` , `Path` , and `Shm` message inputs; unix-socket connections always retain the capability, and the option must not be exposed to untrusted networks. It defaults to `true` in 4.1.5 with a startup warning on every TCP listener, but will default to `false` in the next major release, so TCP clients supplying file paths or shared-memory inputs must migrate to a unix socket before then. The proxy also strips client-supplied `Shm*` headers and `File` , `Path` , `Shm` query arguments, so scanning local files through the proxy no longer works and such clients must connect directly to a scanner listener.

Explicitly opt in to file and shared-memory message inputs on TCP before the default changes to false in the next major release.

```hcl
allow_file_and_shm_inputs = true;
```

Opt in to the future-default now to silence TCP startup warnings on workers that have no clients using file or SHM inputs.

```ini
# local.d/worker-normal.inc
allow_file_and_shm_inputs = false;
```

— Exact config keys across three files with runnable examples and migration deadline.

4.1.5

02 **Structured fuzzy match results in Elastic plugin** BREAKING

90

The `fuzzy_nested` option in `local.d/elastic.conf` switches the fuzzy results field from a flat `fuzzy_hashes` list to the Elasticsearch `nested` type, with per-match objects from task:get_fuzzy_results() containing `rule`, `symbol`, `found`, `queried` hashes, `type`, `probability` / `prob`, `flag`, and `exactness` / `exact`; the schema also adds an explicit mapping for message `size`. Saved searches, dashboards, and ingest pipelines referencing the old flat field must be updated — indices created after the upgrade use the new mapping shape while pre-upgrade indices retain the old one.

Enable structured nested fuzzy match objects in Elasticsearch to get per-match rule, probability, flag, and hash fields instead of a flat list.

```hcl
fuzzy_nested = true;
```

— Exact config key, field names and a runnable example given. 4.1.5

03 **Fuzzy sender facts sharing over encrypted rules** BREAKING

88

Rules sharing extensions with their storage now attach a sender-facts extension to fuzzy check commands sent over encrypted connections, carrying SPF, DKIM, and DMARC results, PTR state (including generic/dynamic hostname verdict from hfilter), bucketed recipient count, and TLS usage; `FUZZY_CALLBACK` gains ordering dependencies so authentication results are available when commands are built. A new `no_share = true` per-rule flag in `local.d/fuzzy_check.conf` disables sender-facts sharing on specific fuzzy rules. Fuzzy storage nodes older than 4.1.5 cannot skip this new extension type and will reject the entire fuzzy command with no reply to scanners — all fuzzy storage nodes must be upgraded to 4.1.5 before scanners, or `no_share = true` set on affected rules.

Temporarily disable sender-facts extension sharing on a fuzzy rule whose storage cannot yet be upgraded to 4.1.5.

```ini
# local.d/fuzzy_check.conf, inside the affected rule
no_share = true;
```

— Names extension fields, callback and config key, plus upgrade order requirement. 4.1.5

04 **Controller and proxy worker connection and auth limits** NEW

86

`max_auth_failures` and `auth_failure_window` in `local.d/worker-controller.inc` rate-limit failed controller password attempts per source (default: 10 failures in 60s triggers HTTP 429; set `max_auth_failures = 0` to disable). `max_connections` and `max_connections_per_source` limits are also added to the controller and proxy workers.

Tune controller auth rate-limiting to allow more attempts before throttling, or disable it entirely for a trusted internal deployment.

```ini
# local.d/worker-controller.inc
max_auth_failures = 20;
auth_failure_window = 120s;
```

— Exact config keys, defaults and a runnable tuning example. product docs

05 **GPT module condition callback signature change** BREAKING

75

The `gpt` module's `condition` callback in `local.d/gpt.conf` now returns only a boolean verdict plus an optional reason string, instead of both a decision and request content; request content generation moves internally into `gpt_check`. Deployments whose condition callback also returned content will have that content silently ignored and must update it to return the boolean decision alone.

— Names config file, callback and exact behavior change, no code sample. 4.1.5

06 **encrypted_only no longer exempts local peers** BREAKING

75

`encrypted_only = true` in `local.d/worker-fuzzy.inc` no longer exempts peers with local or link-local source addresses — all clients, including those on localhost, must now use encryption (`encryption_key` in their `fuzzy_check` rule) or `encrypted_only` must be removed.

— Exact config file, option and required remediation named. product docs

07 **Per-key forbidden_ids no longer merges with worker list**

73

BREAKING

Per-key `forbidden_ids` ACLs in fuzzy storage no longer additively merge with the worker-level `forbidden_ids` option — keys that previously relied on the global list to suppress additional flags must have those flag IDs added explicitly to their own `forbidden_ids` list.

— Names the exact option and before/after behavior, no config path given. product docs

08 **URL filtering and symbol firing fixes** BREAKING 68

url:get_urls_filtered() called without an include list (used by `rspamadm mime urls` and the `URI_COUNT_ODD` rule) no longer drops flag-less URLs such as those extracted from HTML parts, correcting the parity check for multipart/alternative messages. As a result, `HFILTER_HELO_IP_A` and `URI_COUNT_ODD` now fire in cases where they previously did not — local score overrides and composites involving these symbols may produce different scores after upgrade.

— Names the function, command and two symbols affected by the fix. product docs

09 **Font-aware PDF text extraction** NEW 65

Adds font-aware PDF simple font text decoding to libmime, supporting the three Annex D base encodings and per-font `/Differences` overrides, with a builder producing a single UTF-8 buffer per page. Also adds PDF `/ToUnicode` CMap parsing for composite (Type0) fonts, supporting codespace ranges, `bfchar`, `bfrange`, and UTF-16BE destinations including surrogate pairs; the CMap takes precedence over `/Encoding`.

— Detailed mechanism but no user-facing config or flag to act on. 4.1.5

THINNER COVERAGE BELOW

10 **Scanner backpressure from in-flight TCP connections** IMPROVED 43

The scanner's `max_tasks` counter now includes accepted TCP connections still transmitting their body, enabling earlier backpressure on saturated scanners.

— Describes the mechanism but has no config knob or number. product docs

└──▷ **BREAKING ON UPGRADE**

! The `condition` callback for the GPT plugin now returns only a boolean (plus an optional reason) instead of both a decision and request content; deployments whose condition returned request content must update it to return the boolean decision alone.

! The `allow_file_and_shm_inputs` default will become `false` in the next major release; operators relying on file or shared-memory inputs over TCP should set `allow_file_and_shm_inputs = true` explicitly now.

! Fuzzy storage nodes older than 4.1.5 cannot skip the new sender-facts extension type and will reject the entire fuzzy command, returning no reply to scanners — upgrade all fuzzy storage nodes to 4.1.5 before upgrading scanners, or set `no_share = true` on affected rules in `local.d/fuzzy_check.conf`.

! Per-key `forbidden_ids` ACLs in fuzzy storage no longer additively merge with the worker-level `forbidden_ids` option — keys that previously relied on the global list to suppress additional flags must have those flag IDs added explicitly to their own `forbidden_ids` list.

! `encrypted_only = true` in `local.d/worker-fuzzy.inc` no longer exempts peers with local or link-local source addresses — all clients, including those on localhost, must now use

encryption (`encryption_key` in their `fuzzy_check` rule) or `encrypted_only` must be removed.

! The proxy now strips client-supplied `Shm*` headers and `File` , `Path` , `Shm` query arguments — scanning local files through the proxy no longer works; such clients must connect directly to a scanner listener.

! The `allow_file_and_shm_inputs` option defaults to `true` in 4.1.5 (with a startup warning on every TCP listener), but will default to `false` in the next major release — TCP clients supplying file paths or shared-memory inputs must migrate to a unix socket before that release.

! The `gpt` module's `condition` callback in `local.d/gpt.conf` must now return only a boolean and an optional reason string — callbacks that also return content will have that content silently ignored.

! The Elastic module's flat fuzzy field is replaced by structured per-match objects; saved searches, dashboards, and ingest pipelines referencing the old flat field must be updated, and indices created after the upgrade use the new mapping shape while pre-upgrade indices retain the old one.

! `HFILTER_HELO_IP_A` and `URI_COUNT_ODD` now fire in cases where they previously did not — local score overrides and composites involving these symbols may produce different scores after upgrade.

! GPT: The condition callback now returns only a boolean; request content generation has been moved into `gpt_check` .

! Workers: `allow_file_and_shm_inputs` now ships as an explicit `true` default for the normal, controller, and proxy workers — the default will become `false` in the next major release.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ 1 RELEASE · 2026-08-16 CHANGELOG ↗

Sysdig Secure is a container security platform providing runtime monitoring, vulnerability scanning, and compliance enforcement for cloud-native applications.

> Sysdig Secure adds new drill-down views for vulnerability findings—policy evaluation and image layer analysis—plus CIS benchmark automation via CNI plugin detection and base image allow/deny list enforcement in the CLI Scanner.

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

01 Image layer view and filter in Vulnerability Findings NEW

75

Adds a Layers tab to the Resource Details drawer for container image resources, showing the image layer stack grouped by base image or Application Layers, with the Dockerfile instruction and per-severity vulnerability counts per layer. Also adds a Layer filter to the CVE list in Vulnerability Findings so findings can be sliced directly by image layer.

— Names tab, filter and grouping mechanism changelog-20260816-3c434e60

02 CNI plugin detection for CIS benchmark automation NEW

75

KSPM Analyzer v1.49.3 adds detection of Container Network Interface (CNI) plugins in use on each Kubernetes node, enabling automatic evaluation of CIS Kubernetes Benchmark v2.0.0 recommendation 5.3.1 without manual attestation.

— Names version and exact benchmark recommendation, no UI path

changelog-20260816-3c434e60

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Alibaba Cloud OpenSandbox

Feedback ⓘ

1 RELEASE · 2026-08-13 NOTES ↗

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

> OpenSandbox's execd v1.0.22 release adds background job execution with status polling, read-only PTY viewing, SDK session reattachment, and finer-grained isolation-mode probing, while tightening private-network sessions to fail closed at startup.

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

01 **Background runs with status polling** NEW 90

The `run` API now accepts `background: true`, returning a JSON handle so long-running sessions don't block the caller. New `runs/{runId}` and `runs/{runId}/logs` endpoints let callers poll status and stream logs for these background jobs.

Submit a long-running job in the background and poll its status without holding an open connection.

```
$ # Start a background run
curl -X POST https://execd/run \
  -H 'Content-Type: application/json' \
  -d '{"background": true, "command": "python train.py"}'

# Poll status using the returned runId
curl https://execd/runs/<runId>

# Stream logs
curl https://execd/runs/<runId>/logs
```

— Runnable curl example against named endpoints and flag. docker/execd/v1.0.22

02 **SDK session reattachment via attach()** NEW 75

All five SDKs gain attach(sessionId), which rebuilds a fully functional session handle from a bare session ID (e.g. after a worker restart). `SessionState` now echoes the session's original creation parameters.

Reconnect to an existing session from SDK code after a worker restart, without needing the original session object.

```python
session = client.attach(sessionId)
```

— Code example plus named method and field. docker/execd/v1.0.22

03 **Per-mode uid isolation probing** IMPROVED 75

execd now independently probes `setpriv` / `userns` support via `uid_mode` availability checks, returning `503 NOT_SUPPORTED` for unavailable isolation modes instead of failing all isolation globally when one mode is unsupported.

— Named mechanism and status code, no example command. docker/execd/v1.0.22

04 **Private-network sessions fail closed at startup** BREAKING 70

Sessions created with `share_net: false` now bind-pin the workload NetNS/UserNS and return a retryable `503` on failure; previously a bind-pin failure did not block session creation, but now it fails closed at startup.

— Clear before/after and config key, no runnable example. docker/execd/v1.0.22

05 **Read-only PTY viewer mode** NEW 65

A new `?mode=viewer` query parameter on PTY sessions lets multiple read-only clients watch a live PTY alongside the exclusive holder, supporting replay, live output, and backpressure isolation between viewers.

— Named query param but no worked example given. docker/execd/v1.0.22

└──▷ **BREAKING ON UPGRADE**

! `share_net: false` sessions now fail closed at startup: execd must successfully bind-pin the workload NetNS/UserNS or it returns a retryable `503`; previously, failure did not block session creation.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

Feedback ⓘ 1 RELEASE · 2026-08-16 CHANGELOG ↗

Twingate is a zero-trust network access platform that securely connects remote users and devices to private resources without VPN.

> Twingate added HTTP web app protection that injects identity headers into requests, plus a new scoped Helpdesk admin role for support staff.

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

01 **HTTP web app protection with JWT/header injection** NEW

53

Twingate now injects a JWT identifying the user, plus custom headers carrying user and group attributes, into requests sent to protected web apps. This enables integration with both legacy and modern apps without requiring re-authentication at the app layer.

— Explains mechanism but no config keys or endpoints given changelog-20260816-24905ea2

02 **Helpdesk admin role for scoped support access** NEW

52

A new Helpdesk role can be assigned to support staff, granting read-only access to select Admin Console areas plus scoped support actions: resetting MFA, verifying and unverifying devices, and managing trusted serial numbers.

— Names specific permissions but no exact UI path or API changelog-20260816-24905ea2

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ 1 RELEASE · 2026-08-16 CHANGELOG ↗

Firezone is an open-source VPN and secure access platform that replaces traditional corporate VPNs with modern zero-trust networking.

> Firezone's REST API gained a wave of new endpoints for Intune device visibility, resource pool members, and gateway CRUD, plus broad filtering/pagination support across list endpoints, while the desktop/Linux client shipped a headless macOS CLI and doubled Linux tunnel throughput via GRO/GSO.

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

01 **Query filtering and pagination across list endpoints** IMPROVED

80

Adds query params for filtering/pagination to a wide set of list endpoints: `GET /actors` (`email`, `name`, `type`), `GET /clients` (`firezone_id`, `name`), `GET /groups` (`directory_id`, `entity_type`, `name`), `GET /policies` (`group_id`, `resource_id`), `GET /resources` (`address`, `ip_stack`, `name`, `site_id`, `type`), `GET /sites` (`name`), `GET /sites/{site_id}/gateways` (`ipv4`, `ipv6`, `name`), and the identity provider/directory list endpoints `GET /email_otp_auth_providers`, `GET /entra_auth_providers`, `GET /entra_directories`, `GET /google_auth_providers`, `GET /google_directories`, `GET /oidc_auth_providers`, `GET /okta_auth_providers`, and `GET /okta_directories`, each of which also gained `limit` and `page_cursor` params with a corresponding response schema change for cursor-based pagination.

– Names every endpoint and param but not the pagination mechanics 1.0.0

02 **Gateway creation and update endpoints** NEW

71

Adds `POST /sites/{site_id}/gateways` to create a gateway within a site, plus `PATCH /sites/{site_id}/gateways/{id}` and `PUT /sites/{site_id}/gateways/{id}` to update an existing gateway; the corresponding `GET` / `DELETE /sites/{site_id}/gateways/{id}` response schemas also changed to match.

– Full CRUD surface named, exact methods and paths given 1.0.0

03 **Intune device integration endpoints** NEW

65

Adds `GET /intune_devices`, `GET /intune_devices/{id}`, and `GET /intune_integration` endpoints, providing read access to Microsoft Intune-managed device records and integration status via the Firezone REST API.

– Endpoints named but no detail on payload or use case 1.0.0

04 **Resource pool member endpoints** NEW 65

Adds `GET /resources/{resource_id}/pool_members`, `PATCH /resources/{resource_id}/pool_members`, and `PUT /resources/{resource_id}/pool_members` to read and update the pool members associated with a resource.

— Endpoints named but pool-member semantics unexplained `1.0.0`

THINNER COVERAGE BELOW

05 **Linux tunnel throughput via GRO/GSO** IMPROVED 57

Enables GRO and GSO on the TUN device on Linux kernels 6.2 and later, roughly doubling single-stream TCP throughput through the tunnel (1.6.0, #14011).

— Clear mechanism and measured impact, but no user-facing config

`changelog-20260816-c695dd5c`

06 **Schema changes to actors, policies and resources endpoints** 50

BREAKING

Request bodies changed for `PATCH /actors/{id}` and `PUT /actors/{id}`; request and/or response schemas changed for `POST /policies`, `PATCH /policies/{id}`, `PUT /policies/{id}`, `DELETE /policies/{id}`, and `GET /policies/{id}`; and request bodies changed for `POST /resources`, `PATCH /resources/{id}`, and `PUT /resources/{id}`. No field-level detail on what changed is given.

— Endpoints named but nothing on what fields actually changed `1.0.0`

07 **firezone-cli for headless macOS management** NEW 35

Firezone 1.5.19 adds `firezone-cli` for headless macOS management.

— Named tool but no commands, flags or usage shown `changelog-20260816-c695dd5c`

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Nudge Security

Feedback ⓘ 1 RELEASE · 2026-08-16 CHANGELOG ↗

Nudge Security monitors SaaS application usage and security risks across organizations to reduce cloud infrastructure threats.

> Nudge Security shipped a new risk-scoring dashboard, a dedicated passkey adoption dashboard, and an in-browser nudge that intercepts sensitive data before it reaches AI tools, with an optional enforcement mode.

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

01 In-browser AI data-loss nudge with enforcement NEW

80

A new in-browser nudge, configured under `Settings > In-browser nudges`, intercepts sensitive data (PII, secrets, financial info) detected in AI-tool conversations, shows the employee exactly what was found, and offers one-click redaction. An optional enforcement mode blocks sending until the sensitive data is redacted.

Prevent unredacted PII or secrets from reaching AI tools across your org by enabling enforcement mode for in-browser nudges.

📍 In the Nudge Security console, go to Settings › In-browser nudges, enable sensitive data nudges, then turn on enforce redaction so employees cannot send AI conversations containing detected sensitive data without redacting it first.

— Named config path and enforcement toggle with usage steps changelog-20260816-699c2889

02 Passkey adoption dashboard and per-user tracking NEW

73

The new `Dashboards > Passkey usage` dashboard, populated by the browser extension, surfaces passkey adoption rate, top apps using passkeys, a storage provider breakdown toggleable between passkey count and user count, device-loss survivability, and per-login verification status. Passkey usage is also now shown per individual employee on their user details page.

Audit where your organization's passkeys are stored and whether they survive device loss — useful before rolling out a passwordless policy.

📍 In the Nudge Security console, go to Dashboards › Passkey usage. Toggle the storage-provider breakdown between number of passkeys and number of users, then click an individual employee row to verify their passkey is device-independent and that the user was verified at login.

— Exact UI path and toggle named, with a walkthrough example changelog-20260816-699c2889

03 **Risk dashboard with residual risk scoring** NEW 70

The new `Dashboards > Risk` dashboard (open beta) shows residual risk scores, a trend chart filterable to business-critical apps, a business-criticality-vs-residual-risk matrix with drill-down, and a ranked highest-risk-apps table built automatically from Nudge's per-app risk scoring.

— UI path and named components given, no scoring mechanism detail

changelog-20260816-699c2889

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Daytona

Feedback ⓘ 1 RELEASE · SEEN 2026-08-16 NOTES ↗

Daytona is a development environment platform that enables standardized, reproducible coding spaces for teams and AI agents.

> Daytona tightened security in this release with restricted CLI config file permissions and enforced TLS verification in its Ruby Socket.IO client.

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

01 **TLS verification enforced in Ruby Socket.IO client** BREAKING 25

TLS verification is now enforced in the Ruby Socket.IO client.

— States the change but no mechanism or migration detail. snapshot-20260816

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

Feedback ⓘ 1 RELEASE · 2026-08-16 NOTES ↗

CODE ↗

Lightweight coding agent that runs in your terminal

> Codex CLI's alpha channel added disk-space and Windows Dev Drive diagnostics to `codex doctor`, non-blocking environment attachment configuration for threads, and W3C trace-context propagation through exec-server relay frames, alongside smaller TUI, exec-thread, and gRPC health-check improvements.

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

01 **Disk space and Dev Drive checks in codex doctor** NEW

73

`codex doctor` now checks free disk space for both `CODEX_HOME` and the current worktree, showing a warning below a 5 GiB threshold and failing below 1 GiB; on Windows targets it also checks for a Dev Drive.

— Names the exact command and thresholds; run `codex doctor` to see it.

rust-v0.148.0-alpha.20

02 **Non-blocking pending environment attachment for threads** NEW

70

Threads can now start without blocking while environment attachment configuration is still pending, while ensuring turns do not use an attachment prematurely; adds a Failed configuration state and an `environment_failed` callback for per-thread failure and recovery.

— Describes mechanism and names the new state and callback, no usage steps.

rust-v0.148.0-alpha.20

03 **Trace context propagation through exec-server relay** IMPROVED

60

Propagates W3C `traceparent` and `tracestate` trace context through exec-server relay frames, including fragmented encrypted Noise requests.

— Names exact headers and mechanism but is internal, not directly actionable.

rust-v0.148.0-alpha.20

THINNER COVERAGE BELOW

04 **Health endpoint on code-mode gRPC listener** NEW

40

Adds a health endpoint to the code-mode gRPC listener.

— Names the exact surface added but no path or usage example given.

rust-v0.148.0-alpha.20

05 **TUI startup and context state improvements** IMPROVED

35

TUI app directory state is now scoped to the active context, and the TUI shows resume and fork status during startup.

— Thin descriptions of two small TUI behavior changes. rust-v0.148.0-alpha.20

06 **Workload identity context in token exchange** IMPROVED

35

Forwards workload identity context during token exchange.

— Names the mechanism briefly but lacks further detail. rust-v0.148.0-alpha.20

07 **Paginated history for persistent exec threads** IMPROVED

25

Persistent exec threads now use paginated history.

— One-line description with no mechanism or usage detail. rust-v0.148.0-alpha.20

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

vMLX

Feedback ⓘ 1 RELEASE · 2026-08-15 NOTES ↗

vMLX - JANGTQ Uber Compressed MLX Models - L2 Disk Cache (survives restart) + L1 Paged (super fast ttft) + Hybrid SSM Scheduler + Cont Batching + etc!

> vMLX 1.6.29 focuses on cache reuse across long-context and multiturn workloads, adding a hybrid prefix cache for much faster follow-up turns and a default-on chunked SSM re-derive path that restores cache reuse at scale, alongside verification across a dozen model families.

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

01 **Chunked SSM re-derive and long-context cache stability**

IMPROVED 73

Chunked SSM re-derive ships enabled by default, restoring long-context cache reuse from 0% to 99.9% above 12.5k tokens on hybrid SSM model families. This was validated with a 100k-context multiturn test where a 97.6k-token conversation retrieved planted facts with 99.9% cache reuse and no memory faults.

— Detailed metrics and default behavior, but no user-facing control named v1.6.29

02 **Hybrid prefix cache for long-document follow-ups** NEW

63

A hybrid prefix cache delivers up to 20x faster follow-up turns on long documents: a 43.7k-token document follow-up drops from 107s to 5s, with output byte-identical at temperature 0.

— Concrete before/after numbers but no flag or config to invoke it v1.6.29

THINNER COVERAGE BELOW

03 **Per-family cache and generation hardening** IMPROVED

50

Cache, parser, and generation-default hardening was verified across 12 model families, including DSV4 Flash native composite caching, Gemma 4 mixed-SWA with audio, Qwen 3.6/3.8 hybrid lines with native MTP, and TurboQuant KV bundles.

— Names families tested but no actionable steps or config given v1.6.29

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Certiv

Feedback ⓘ 1 RELEASE · SEEN 2026-08-16 NOTES ↗

Certiv is an AI-powered code security platform that identifies and remediates vulnerabilities in software applications.

> Certiv shipped a new Privacy Mode that ties agent session visibility to policy risk with auditable access logging, alongside its first published API for lead-capture endpoints.

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

01 **Privacy Mode for agent session visibility** NEW 72

Privacy Mode changes session visibility defaults so only sessions with policy violations are automatically surfaced to security teams, while compliant sessions stay private unless an authorized user records an explicit justification for access. Certiv logs who accessed a private compliant session and the recorded reason, creating an auditable trail for governance, privacy teams, and works councils. It integrates with the existing Policy Engine and Scout endpoint observer, so policy verdict determines session visibility without altering real-time Block, Pause, and Steer interventions. Privacy Mode is now available in Preview to existing Certiv customers, enabled on request through the Certiv team.

– Mechanism and integration points clear, but enablement is a support request, not a config key.

Enterprise Security with Privacy by Des…

THINNER COVERAGE BELOW

02 **Public API with lead capture endpoints** NEW 38

Certiv now publishes an API covering one area, Lead Capture, with 2 public endpoints for requesting a demo or subscribing to the newsletter.

– Endpoint count and purpose given but no methods or paths named. 1.0.0

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ 1 RELEASE · 2026-08-15

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 v2.9.5 adds support for three new inference providers, opt-in baseline auto-discovery, and static detection for insecure deserialization, alongside two YARA matching accuracy fixes.

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

01 **Auto-discovery of baseline config file** NEW 60

Adds opt-in auto-discovery of a top-level `.skillspector-baseline.yaml` file, keeping explicitly supplied baselines authoritative when both are present.

— Names exact config file and precedence behavior v2.9.5

THINNER COVERAGE BELOW

02 **Static detection for insecure deserialization** NEW 50

Adds static analyzer coverage for insecure deserialization patterns represented by findings `AST10`, `TT6`, `DS1`, `DS2`, `DS3`, and `DS4`.

— Lists specific finding codes but no detection mechanism detail v2.9.5

03 **Ollama, Azure OpenAI, and generic OpenAI-compatible providers** NEW 45

Adds Ollama support for local OpenAI-compatible inference, Azure OpenAI deployment routing, and a configurable provider for generic OpenAI-compatible endpoints.

— Names three providers but no config keys or setup steps v2.9.5

04 **YARA post-filter rule namespace scoping** IMPROVED 45

Scopes the destructive-autonomy YARA post-filter to SkillSpector's built-in rule namespace and preserves deterministic built-in rule precedence, preventing custom YARA rules that reuse a built-in rule name from being incorrectly post-filtered.

— Clear mechanism and fixed bug scenario, no direct user action v2.9.5

05 **Byte-offset YARA match line mapping** IMPROVED 40

Switches YARA match-to-source-line mapping to byte offsets so non-ASCII content is reported at accurate locations.

— Explains before/after but no user-facing action v2.9.5

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 16, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue