THE CYBER TOOLCHAIN

NO. 081

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 21, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 189 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 24 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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

Three unrelated things stand out. Trend Vision One's local response filters are GA, so a matching process gets killed on the Windows

vim ? help

endpoint itself instead of waiting on a cloud round-trip. BunkerWeb adds mutual TLS to upstreams and a CrowdSec investigation API. NETRESEC's FlowCarp fingerprints C2 families from flow statistics, catching beacons that carry no signature.

RESPOND

Terminate a malicious process on the endpoint without a cloud round-trip

Local response filters are now GA and auto-terminate matching processes on Windows endpoints. Containment executes on the host rather than after a detection travels to the console and a task comes back, which is the difference between stopping a loader and cleaning up after it.

Trend Vision One

DETECT

Catch C2 that has no signature, and hunt OT protocols in PCAP

FlowCarp classifies C2 protocols from packet behaviour rather than byte patterns, so a family that changes its payload or moves to a fresh domain still looks the same statistically; the CLI trains custom models on your own captures. Netcap adds structured Modbus and S7comm parsing with CISA-mapped rules, so ICS traffic can be hunted as fields instead of eyeballed in Wireshark.

NETRESEC · Netcap

PROTECT

Authenticate the reverse proxy to its own backends

Mutual TLS to upstreams means the origin can now verify the proxy is the proxy, closing the gap where anything that reaches the backend network is trusted. The new CrowdSec investigation API and UI also let an analyst pull the decision history behind a block instead of guessing why a request was dropped.

BunkerWeb

PROTECT

Block a bad NuGet package before it lands, and see C/C++ dependencies at all

Package Firewall now covers NuGet, so a known-malicious or policy-violating .NET package is refused at fetch time rather than found in a scan

afterwards. Conan support brings C/C++ into composition analysis, a stack that traditionally has no lockfile to read and so goes uninventoried.

Endor Labs

IDENTIFY

Send an externally discovered web app straight into a scanner

EASM findings can be pushed directly into DAST Essentials, removing the copy-paste step between "we found an unknown internet-facing app" and "we tested it." Attack-surface discovery only pays off when the queue it produces gets worked, and this shortens that handoff.

Veracode

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Keygraph shannon

Feedback ⓘ v3.3.0 NOTES ↗ RANK

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 v3.3.0 adds live refresh of its AI model catalogue at scan start, so new models become available without requiring a tool update.

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

01 **Network refresh of AI model catalogue** IMPROVED 28

Shannon now refreshes its model catalogue over the network at scan start, ensuring newly available AI models are accessible without updating the tool.

— Describes behaviour but no mechanism, config, or command detail. v3.3.0

v3.3.0 NOTES ↗ LATEST STABLE

Shannon v3.3.0 refreshes the AI model catalogue over the network at scan start so new models are available without a tool update.

└──▷ **GET THIS VERSION**

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

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

cve-lite-cli

Feedback ⓘ v1.36.0 NOTES ↗ RANK

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 v1.36.0 opens up its PR delta renderer for external use and patches two false-result bugs from the prior release.

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

01 **Public PR delta renderer package** NEW 50

The PR delta renderer, which generates pull request comments by diffing two scan reports, is now published as part of the distributed cve-lite-cli package. This lets external tools and CI workflows outside the repository generate PR comments from scan diffs without reimplementing the logic.

– Names the module and its function but no CLI flag or command v1.36.0

02 **False-result bug fixes** IMPROVED 10

Fixes two false-result bugs that were introduced in v1.35.0; the specific bugs are not described.

– No details on which false results or how fixed v1.36.0

v1.36.0 NOTES ↗ LATEST STABLE

cve-lite-cli v1.36.0 ships the PR delta renderer as a public package and fixes two false-result bugs from v1.35.0.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.36.0 https://github.com/OWASP/cve-lite-cli.git
# already have the repo? check out this version:
$ git checkout v1.36.0
```

Publishes the PR delta renderer as part of the distributed package, enabling external tools and CI workflows outside this repository to generate pull request comments from two scan reports.

WAS THIS USEFUL?  

Aikido Security

Feedback ⓘ announced · Sep 21, 2026 BLOG ↗

RANK

Aikido Security is a developer-focused security platform that identifies and remediates vulnerabilities in code and dependencies.

> Aikido released Altar, an open-weight AI security model built to run Aikido Machine's autonomous pentesting entirely inside customer-controlled or air-gapped infrastructure.

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

01 **Altar open-weight AI security model** NEW 83

Altar is an open-weight security AI model (328 GB, `W4A16` quantization, 168 of 256 routed experts retained) that powers Aikido Machine's autonomous pentesting entirely inside customer-controlled or air-gapped infrastructure, without sending code or findings to a third-party inference service. It is deployable on a 4-H200s node using `vLLM`, with model weights, a model card, license, serving flags, and deployment instructions published under Aikido's organization.

.
.
.
.

— Names exact specs and deployment path, but no runnable command shown.

launch-20260921-b3ced206

> announced · Sep 21, 2026 BLOG ↗
>
> Aikido releases Altar, an open-weight AI security model for air-gapped and sovereign deployments inside Aikido Machine.
>
> Introduces Altar, an open-weight security AI model (328 GB, W4A16 quantization, 168 of 256 routed experts retained) that powers Aikido Machine's autonomous pentesting entirely inside customer-controlled or air-gapped infrastructure without sending code or findings to a third-party inference service.
>
> Altar is deployable on a 4-H200s node using `vLLM`, with model weights, a model card, license, serving flags, and deployment instructions available under Aikido's organization.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ EASM updates (2026-09-21) BLOG ↗ RANK

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities in code.

> Veracode's EASM module gets a full UI redesign, a new Findings view, Dashboard widgets, Business Units management, and a direct path to push discovered web apps into DAST Essentials for scanning, alongside real-time scan notifications, multi-project visibility, scan scheduling and AI-driven risk prioritization.

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

01 Export EASM assets to DAST Essentials as scan targets NEW

75

Enables export of web applications from EASM directly into DAST Essentials by marking assets as DAST Candidates, which creates discovered targets in DAST Essentials ready for authenticated scanning.

Promote a discovered web application to a DAST scan target without leaving EASM — useful when an asset surfaces during footprint discovery and you want to immediately queue it for authenticated scanning.

> 📍 In the EASM UI, open the asset you want to scan, mark it as a DAST Candidate, then navigate to DAST Essentials to find the asset listed as a discovered target ready for scanning.

— Named mechanism and step-by-step usage example provided EASM updates (2026-09-21)

THINNER COVERAGE BELOW

02 Dashboard widgets: Heatmap, World map, Things NEW

50

Adds Dashboard widgets — Heatmap, World map, and Things — with filterable data views and a shortcut to create reports.

— Names the three widgets but not the filter options EASM updates (2026-09-21)

03 Real-time notifications, visibility, scheduling and AI prioritization in EASM

NEW

45

Adds real-time in-app notifications when scans complete, fail, or are cancelled; multi-project visibility for centralized risk management across all projects; scan scheduling to streamline orchestration; AI-powered intelligence for accelerated risk prioritization;

and an Issues capability to streamline remediation efforts and improve focus on cyber hygiene.

— Five thin one-line additions grouped, none with a named UI path or config

EASM updates (2026-09-21)

04 **Findings view in EASM** NEW 45

Adds a Findings view to surface large volumes of security data as actionable insights, with prioritization workflows and expanded reporting.

— Describes purpose but not exact fields or filters EASM updates (2026-09-21)

05 **Redesigned EASM UI with Legacy UI option** IMPROVED 40

Introduces a redesigned EASM UI with improved navigation and usability. The Legacy UI remains available for users who aren't ready to switch.

— Names the change but no navigation path or scope detail EASM updates (2026-09-21)

06 **Business Units management in EASM** NEW 35

Adds Business Units management, enabling creation, assignment, and domain association.

— States capability without mechanism or UI location EASM updates (2026-09-21)

EASM updates (2026-09-21) BLOG ↗ LATEST STABLE

Veracode EASM gains redesigned UI, real-time scan notifications, multi-project visibility, scan scheduling, and AI-powered intelligence.

└──▷ **HOW TO FIND IT**

Promote a discovered web application to a DAST scan target without leaving EASM — useful when an asset surfaces during footprint discovery and you want to immediately queue it for authenticated scanning.

📍 In the EASM UI, open the asset you want to scan, mark it as a DAST Candidate, then navigate to DAST Essentials to find the asset listed as a discovered target ready for scanning.

Adds real-time in-app notifications when scans complete, fail, or are cancelled.

Adds multi-project visibility for centralized risk management across all projects.

Adds scan scheduling to streamline scan orchestration across projects.

Adds AI-powered intelligence for accelerated risk prioritization.

Introduces redesigned EASM UI with improved navigation and usability, with the Legacy UI still available.

+5 more

WAS THIS USEFUL?  

◆ Secrets Detection & Management

infisical

Feedback ⓘ v0.165.14 NOTES ↗ RANK

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

> Infisical's latest release adds PowerDNS as a DNS-01 provider for automated certificate issuance, introduces fine-grained request routing controls for Agent Vault, and adds event-triggered alerts alongside billing and secrets-UI improvements.

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

01 **Fine-grained routing rules for Agent Vault** NEW 52

Adds method and path rules, custom headers, and substitutions to Agent Vault for fine-grained request routing and transformation control.

– Lists concrete rule types but no exact syntax or flags. v0.165.14

02 **PowerDNS DNS-01 provider for ACME** NEW 48

Adds PowerDNS as a DNS-01 provider for ACME certificate authorities, enabling automated certificate issuance via PowerDNS-managed zones.

– Names provider and mechanism but no config steps. v0.165.14

03 **Event outbox and event-triggered alerts** NEW 30

Adds a transactional event outbox and event-triggered alerts for automated notifications on secret or infrastructure events.

– Describes capability but no event types or config named. v0.165.14

04 **Bulk secret selection and clipboard copy** IMPROVED 26

Simplifies bulk secret selection and adds clipboard copying to the secrets UI.

– UI change described but no navigation path given. v0.165.14

05 **Billing v2 usage breakdown and org listing** NEW 23

Adds billing v2 usage breakdown and organization listing via a new API implementation.

– No endpoint, fields, or usage details given. v0.165.14

> v0.165.14 NOTES ↗ STABLE
>
> Infisical v0.165.14 adds PowerDNS DNS-01 support, Agent Vault routing rules, and event-triggered alerts

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.165.14
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.165.14
```

Adds PowerDNS as a DNS-01 provider for ACME certificate authorities, enabling automated certificate issuance via PowerDNS-managed zones.

Adds method and path rules, custom headers, and substitutions to Agent Vault for fine-grained request routing and transformation control.

Adds a transactional event outbox and event-triggered alerts for automated notifications on secret or infrastructure events.

Adds billing v2 usage breakdown and organization listing via a new API implementation.

Simplifies bulk secret selection and adds clipboard copying to the secrets UI.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Bunkerity BunkerWeb

Feedback ⓘ v1.6.15 NOTES ↗ RANK

Open-source and cloud-native Web Application Firewall (WAF)

> BunkerWeb 1.6.15 ships a guided, flag-rich Docker upgrade script with automatic backup, health checks and rollback; a full CrowdSec investigation API and UI; mutual TLS to upstreams; and a bundle of new file/IP-reputation plugins (SentinelOne, SysWarden, ClamAV, VirusTotal), alongside Redis-resilient session storage and per-setting draft mode in the Web UI RAW editor.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 Guided Docker upgrade script for BunkerWeb stacks NEW

95

`install-bunkerweb.sh --docker` gains `--compose-dir PATH` (stack directory instead of cwd), `-v, --version VERSION` and `--image-tag TAG` to select the target image, `--backup-dir PATH` for the pre-upgrade backup archive, `--overwrite-compose` to force-regenerate `docker-compose.yml`, `--force-type-change` for destructive topology switches, `--no-pull` to skip image pulls, and `--docker-wait-timeout N` (default 600s) for health-check patience. Before changes it runs `bwcli plugin backup save` in the Scheduler container and backs up `.env` and `docker-compose.yml` as a `.bak.<timestamp>` copy, aborting on backup failure unless overridden; it reads the running version from the live container to reliably handle rolling tags (`latest`, `testing`) and interrupted upgrades; it rewrites `.env` preserving manually added entries and only regenerates `docker-compose.yml` if it still matches the script-generated original; it auto-detects install type, secrets, host ports, worker list and Compose project name from `.env` to prevent accidental credential rotation or port changes; it then runs `docker compose pull` and `docker compose up -d`, recreating only containers whose image changed, requires two healthy observations at least ten seconds apart, exits with status 2 (no rollback) on timeout, backs up the database, and refuses downgrades.

Upgrade while preserving local `docker-compose.yml` edits but skip pulling images (useful when images are pre-pulled in an air-gapped environment).

```
$ bash install-bunkerweb.sh --docker --compose-dir /opt/bunkerweb
-v 1.6.15 --no-pull
```

Force-regenerate `docker-compose.yml` when switching stack topology (e.g. from all-in-one to split manager/worker), accepting the destructive topology change.

```
$ bash install-bunkerweb.sh --docker --compose-dir /opt/bunkerweb
-v 1.6.15 --overwrite-compose --force-type-change
```

Upgrade a stack that lives outside the current directory to a specific version, with a custom backup destination and an extended health-check timeout for slow first starts.

```
$ bash install-bunkerweb.sh --docker --compose-dir /opt/bunkerweb
--version 1.6.15 --backup-dir /mnt/backups --docker-wait-timeout
900
```

Force regeneration of a locally modified docker-compose.yml when the local edits are no longer needed and you want the script-generated layout back.

```
$ bash install-bunkerweb.sh --docker --compose-dir /opt/bunkerweb
--version 1.6.15 --overwrite-compose
```

— Every flag named with runnable upgrade commands and rollback mechanics. v1.6.15

---

02 **CrowdSec investigation and management API/UI** NEW 95

Adds `CROWDSEC_MANAGEMENT_LOGIN` and `CROWDSEC_MANAGEMENT_PASSWORD` config keys for a dedicated CrowdSec machine identity, plus endpoints `GET /crowdsec` (connection status/sync/errors), `GET /crowdsec/{connection_id}/decisions` (filterable by `origin` and `scenario`, paginated via `offset` / `limit`, max 200), `GET /crowdsec/{connection_id}/ips/{ip}` (up to 200 decisions, 50 alerts, 50 reports with totals), `GET /crowdsec/{connection_id}/alerts/{alert_id}`, `GET /crowdsec/{connection_id}/allowlists` (paginated, up to 200 entries per list), `GET /crowdsec/{connection_id}/allowlists/check?ip={ip}`, and `DELETE /crowdsec/{connection_id}/decisions/{decision_id}` (requires `scope`, `value`, `decision_type` in the JSON body). New delegated permissions `crowdsec_read` and `crowdsec_delete` scope this under the `bans` resource, and a new Additional Pages → CrowdSec UI page consolidates active decisions, alerts, BunkerWeb reports, local bans and native allowlist status with pre-filled IP lookup links.

Investigate all active CrowdSec decisions for a specific IP across a connection, to confirm whether a block is still in force before escalating a support ticket.

```
$ curl -s -H 'Authorization: Bearer <api_token>'
'https://bunkerweb.example.com/api/crowdsec/<connection_id>/ips/2
03.0.113.42'
```

Remove a mistaken CrowdSec ban on an IP so it propagates to all bouncers at their next stream update.

```
$ curl -s -X DELETE -H 'Authorization: Bearer <api_token>' -H
'Content-Type: application/json' -d
'{"scope":"ip","value":"203.0.113.42","decision_type":"ban"}'
'https://bunkerweb.example.com/api/crowdsec/<connection_id>/decis
ions/<decision_id>'
```

Remove an active CrowdSec ban from the local API without touching BunkerWeb-local bans — useful when a decision was issued in error on a shared engine.

```
$ curl -X DELETE
https://bunkerweb.example.com/crowdsec/<connection_id>/decisions/
<decision_id> \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"scope": "ip", "value": "203.0.113.42", "decision_type":
"ban"}'
```

Investigate all active CrowdSec decisions for a suspicious IP, including matching alerts and retained BunkerWeb reports, to triage an incident.

```
$ curl -s
'https://bunkerweb.example.com/crowdsec/<connection_id>/ips/203.0
.113.42' \
  -H 'Authorization: Bearer <api_token>' | jq .
```

Programmatically remove a specific CrowdSec ban on an IP address via the public API, useful in an automated remediation workflow.

```
$ curl -s -X DELETE
https://bunkerweb.example.com/api/crowdsec/<connection_id>/decisi
ons/<decision_id> \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"scope": "ip", "value": "198.51.100.42", "decision_type":
"ban"}'
```

Investigate a suspicious IP across all CrowdSec decisions, alerts, and retained BunkerWeb reports in one call.

```
$ curl -s
'https://bunkerweb.example.com/api/crowdsec/<connection_id>/ips/1
98.51.100.42' \
  -H 'Authorization: Bearer <api_token>'
```

— Every endpoint, method and required field named with runnable curl calls. v1.6.15

03 **Session storage, Redis fallback and revocation controls** NEW

75

`UI_USE_REDIS=no` decouples the Web UI from Redis independently of the global `USE_REDIS` setting; when Redis is unavailable or rejects writes due to `maxmemory`, UI sessions fall back to a local per-host cache, with logout and session-ID rotation revocations recorded in both stores. Destroyed sessions are rejected on next use, locally or cluster-wide when Redis is enabled, and the new `SESSIONS_REVOCATION_MEMORY_SIZE` config key controls the in-memory store size for revoked cookie session identifiers.

Disable Redis for the BunkerWeb UI session store only, keeping Redis active for inter-instance lock and report sharing.

```ini
UI_USE_REDIS=no
```

Isolate the UI from Redis session storage without disabling Redis-backed bannishments and inter-instance reporting — useful when you want Redis for data sharing but local session handling.

```ini
UI_USE_REDIS=no
```

Isolate the UI's Redis dependency from the rest of BunkerWeb — useful when you want bans and reports shared via Redis but need the UI to run without it.

```env
UI_USE_REDIS=no
```

— Named config keys and fallback mechanism, no numeric size default given. v1.6.15

---

04 **Header-based bypass and match rules across security plugins**

NEW

75

Adds numbered header-match config pairs `ANTIBOT_IGNORE_HEADER_NAME` / `ANTIBOT_IGNORE_HEADER_VALUE`, `BLACKLIST_HEADER_NAME` / `BLACKLIST_HEADER_VALUE`, `BLACKLIST_IGNORE_HEADER_NAME` / `BLACKLIST_IGNORE_HEADER_VALUE`, and `COUNTRY_IGNORE_HEADER_NAME` / `COUNTRY_IGNORE_HEADER_VALUE` (with PCRE-matched values) to bypass or trigger the antibot, blacklist and country plugins; header-based matching is also extended to the `whitelist`, `greylist` and `dnsbl` plugins, with header values stored as password settings that are never logged or cached.

Allow an internal health-check service to bypass the antibot challenge by presenting a shared secret header.

```env
ANTIBOT_IGNORE_HEADER_NAME_1=X-Internal-Token
ANTIBOT_IGNORE_HEADER_VALUE_1=^my-secret-token$
```

Allow internal monitoring probes to skip the antibot challenge by presenting a shared secret header, without opening access to all traffic.

```ini
ANTIBOT_IGNORE_HEADER_NAME_1=X-Internal-Token
ANTIBOT_IGNORE_HEADER_VALUE_1=^my-secret-rotation-value$
```

Allow a trusted internal service to bypass the antibot challenge by presenting a shared-secret header, while anchoring the regex so partial matches are rejected.

```env
ANTIBOT_IGNORE_HEADER_NAME_1=X-Internal-Token
ANTIBOT_IGNORE_HEADER_VALUE_1=^my-secret-token$
```

— All header config keys enumerated with a working config example. v1.6.15

05 **Per-setting draft mode in RAW editor** NEW 65

In the RAW editor, prefix a `KEY=value` line with `#` to mark it draft (value stored but not applied, shown with a muted `#` prefix, inherited/default value takes effect); press `Backspace` on a draft line to reactivate it. Draft lines persist across saves in RAW mode and are preserved when saving in Simple or Advanced mode.

Stage a config change in the RAW editor without applying it yet — useful for preparing a value to activate in a single future save.

> 📍 # In the BunkerWeb web UI RAW editor: # 1. Open a service or global settings RAW editor. # 2. Place the cursor at the start of a `KEY=value` line. # 3. Press # to toggle the line into draft state (line becomes highlighted). # 4. Save — the value is stored but the effective config remains unchanged. # 5. Press Backspace on the draft line (or # again) to activate, then save to apply.

Prepare a config change in the RAW editor without applying it yet — stage it as a draft and activate it in a later save when ready.

> 📍 In the RAW editor, place the cursor at the start of a `KEY=value` line and press # to toggle it to draft state. The line is highlighted and the value is stored but not applied. Press Backspace on the drafted line to activate it, then save to apply.

Temporarily disable a setting in RAW mode without deleting it, so its inherited or default value applies while you keep the line for later activation.

> 📍 In the RAW mode editor, navigate to the start of a setting line and press `#` to mark it as draft. The line gains a muted `#` prefix and the inherited/default value takes effect. Press `Backspace` at the start of the line to reactivate it, then click Save to persist.

— Exact keystrokes given as UI steps, no config-file equivalent. v1.6.15

---

06 **New plugin bundle: SentinelOne, SysWarden, ClamAV, VirusTotal**

NEW 60

`EXTERNAL_PLUGIN_URLS` installs the v1.12 plugin bundle from a URL such as `https://github.com/bunkerity/bunkerweb-plugins/archive/refs/tags/v1.12.zip`. `bunkerweb-plugins/sentinelone` checks uploaded file hashes and client IPs against the SentinelOne API, blocking on reputation or IOCs; `bunkerweb-plugins/syswarden` syncs BunkerWeb bans to SysWarden for kernel-level nftables blocking and imports its block/allow lists; a ClamAV plugin scans uploaded files and rejects requests when malicious; a VirusTotal plugin checks uploaded file hashes and client IPs against the VirusTotal API, blocking when configured reputation thresholds are exceeded.

Install the full v1.12 plugin bundle — including SentinelOne, SysWarden, ClamAV, and VirusTotal — by pointing the scheduler at the tagged archive.

```text
EXTERNAL_PLUGIN_URLS=https://github.com/bunkerity/bunkerweb-plugins/archive/refs/tags/v1.12.zip
```

— Plugins and install key named, each mechanism only one line. product docs

07 **ModSecurity audit log path and concurrent logging** NEW

60

Adds `MODSECURITY_SEC_AUDIT_LOG` config key to configure the ModSecurity audit log path (under `/var/log/bunkerweb` by default), and adds opt-in concurrent audit logging for ModSecurity with instance-side storage validation.

Send ModSecurity concurrent audit log entries to a dedicated path for ingestion by a SIEM, rather than the default log location.

```ini
MODSECURITY_SEC_AUDIT_LOG=/var/log/bunkerweb/modsec_audit.log
```

— Named config key with a concrete example path, but validation detail is thin. v1.6.15

THINNER COVERAGE BELOW

08 **Mutual TLS to upstreams in reverseproxy and grpc plugins** NEW

50

Supports presenting a client certificate to upstreams for mutual TLS in the `reverseproxy` and `grpc` plugins, with configurable upstream CRL, TLS protocols, and ciphers validated and distributed by the Scheduler.

— Mechanism named but no specific config key given. v1.6.15

09 **gRPC plugin feature parity with reverseproxy** IMPROVED

50

Expands the `grpc` plugin to match `reverseproxy` configurability, including client response headers, pass/ignore headers, underscores in headers, buffer and body-size controls, `auth_request`, and upstream TLS verification.

— Feature list named but no exact config keys given. v1.6.15

10 **Per-URL request body size limit** NEW

50

Adds per-URL `REVERSE_PROXY_MAX_CLIENT_SIZE` config key to control both the request body size and the ModSecurity body-size limit on a per-location basis.

– Named config key but no example or default value given. v1.6.15

11 **Two-factor auth anti-replay moved to shared database** BREAKING

50

Replay protection for two-factor authentication codes is now stored in the shared database across Web UI replicas instead of per-replica memory. During a rolling upgrade, a TOTP code already displayed may be rejected for up to 33 seconds; upgrade all Web UI replicas simultaneously to minimise the window, and TOTP encryption keys must remain available to every replica as before.

– Migration guidance given but no config key or command. product docs

12 **logrotate integration for all-in-one image** NEW

45

Adds logrotate integration for the all-in-one image; mount `/etc/logrotate.d/bunkerweb` to provide a custom rotation policy.

– Mount path given but no example policy or defaults. v1.6.15

└──▷ **ALSO FROM THESE RELEASES**

Disable Redis for the Web UI only — keeping global ban and report sharing intact — when the UI workers are hitting Redis connection limits.

```ini
UI_USE_REDIS=no
```

└──▷ **BREAKING ON UPGRADE**

! Two-factor authentication anti-replay state is now stored in the shared database instead of per-replica memory; during a rolling upgrade, a TOTP code already displayed may be rejected for up to 33 seconds — upgrade all web UI replicas simultaneously to minimise the window. TOTP encryption keys must remain available to every replica as before.

! Two-factor authentication replay protection is now stored in the shared database and only covers replicas running this version; during a mixed-version migration window an already-displayed TOTP code may be rejected for up to 33 seconds.

v1.6.15 NOTES ↗ LATEST STABLE

BunkerWeb 1.6.15 adds mTLS to upstreams, per-header allow/block criteria, per-URL client size limits, and a syslog-ng UI log collector.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.6.15
https://github.com/bunkerity/bunkerweb.git
# already have the repo? check out this version:
$ git checkout v1.6.15
```

└──▷ **USE IT**

Disable Redis for the Web UI only — keeping global ban and report sharing intact — when the UI workers are hitting Redis connection limits.

/etc/bunkerweb/variables.env

```
UI_USE_REDIS=no
```

Send ModSecurity concurrent audit log entries to a dedicated path for ingestion by a SIEM, rather than the default log location.

/etc/bunkerweb/variables.env

```
MODSECURITY_SEC_AUDIT_LOG=/var/log/bunkerweb/modsec_audit.log
```

Adds `UI_USE_REDIS=no` to disable Redis exclusively for the Web UI without disabling the global Redis-backed ban and report sharing.

Adds `MODSECURITY_SEC_AUDIT_LOG` to configure the ModSecurity audit log path under `/var/log/bunkerweb`.

Adds per-URL `REVERSE_PROXY_MAX_CLIENT_SIZE` to control both request body size and the ModSecurity body-size limit on a per-location basis.

Supports presenting a client certificate to upstreams for mutual TLS in the `reverseproxy` and `grpc` plugins, with configurable upstream CRL, TLS protocols, and ciphers validated and distributed by the Scheduler.

Adds header-based matching to `antibot`, `blacklist`, `whitelist`, `greylist`, `dnsbl`, and `country` plugins — match on a request header name and optional PCRE value regex as an ignore or list criterion; header values are stored as password settings and are never logged or cached.

+7 more

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite

Feedback ⓘ

Professional / Community 2026.9 (2026-09-21) BLOG ↗ RANK

We enable the world to secure the web. Burp Suite, the Web Security Academy, and world-leading research from PortSwigger â the AppSec communityâs home.

> Burp Suite 2026.9 adds the ability to export SMTP Collaborator interactions as email files, alongside expanded hotkey and title bar customization and command palette performance improvements.

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

01 **SMTP Collaborator interaction export** NEW 35

Adds the ability to export SMTP Collaborator interactions as email files.

– Names the feature but no export mechanism or location detail

Professional / Community 2026.9 (2026-0…

02 **Hotkey and title bar customization** IMPROVED 21

Expands hotkey and title bar customization options.

– Vague on which options were added Professional / Community 2026.9 (2026-0…

03 **Command palette performance improvements** IMPROVED 13

Improves command palette performance.

– No detail on scope or magnitude of improvement Professional / Community 2026.9 (2026-0…

> Professional / Community 2026.9 (2026-09-21) BLOG ↗ LATEST STABLE
>
> Burp Suite 2026.9 adds SMTP Collaborator email export and expands hotkey and title bar customization.
>
> Adds the ability to export SMTP Collaborator interactions as email files.
>
> Expands hotkey and title bar customization options.
>
> Improves command palette performance.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-09-09 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds agent_worktree.py for isolated PR/branch checkouts and SHA256 deduplication in ResultServer

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-09-09
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-09
```

└──▷ TRY IT

Check out a contributor's PR into an isolated worktree so CI bots or reviewers can test it without risking uncommitted work in the main clone.

```
$ python utils/agent_worktree.py new --pr 3219
```

List all managed worktrees with their status flags (main / managed / dirty / pr#N), then tear down all managed ones after testing is complete.

```
$ python utils/agent_worktree.py list
python utils/agent_worktree.py cleanup --all
```

Adds `utils/agent_worktree.py` with subcommands `new --pr`, `new --branch`, `new --from`, `list`, `path`, `update`, `info`, `remove`, and `cleanup --all` for managing disposable `git worktree` checkouts of PRs or branches without touching the working clone; every subcommand accepts `--repo PATH` and `--json` flags.

ResultServer now performs SHA256 deduplication and file versioning instead of raising `CuckooOperationalError` / `FileExistsError` when an uploaded filename already exists in `storage/analyses/<task>/files/`.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

pwndbg

Feedback ⓘ 2026.09.15 NOTES ↗ CODE ↗ RANK

Exploit Development and Reverse Engineering with GDB & LLDB Made Easy

> pwndbg's 2026.09.15 release adds a new stack-vis command, extends nearpc to disassemble whole functions without debug symbols, introduces backwards disassembly at breakpoints, expands page-info commands, enables x64 kernel emulation, and improves handling of segfaulting indirect jumps — alongside a breaking rename in the debugger error API.

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

01 Symbol-free function disassembly in nearpc IMPROVED

80

Expands `nearpc -f` / `--function` to work without debugging symbols, making it a full replacement for GDB's/LLDB's `disass` command with branch visualization, branch targets, and operand annotations.

Disassemble an entire function by address without debug symbols, getting branch visualization and operand annotations that raw `disass` lacks.

```
$ nearpc --function 0x401080
```

▶ View command output 2026.09.15

Real pwndbg output from a synthetic AArch64 core fixture.

— Names flags and comparison target, runnable example given 2026.09.15

02 New stack-vis command for stack frames NEW

70

Adds the `stack-vis` command to visualize stack frames of the current thread, similar to how `vis-heap-chunks` visualizes heap chunks.

Visualize the current thread's stack frames at a breakpoint to understand frame layout during exploit development.

```
$ stack-vis
```

▶ View command output 2026.09.15

Real pwndbg output from a synthetic AArch64 core fixture.

— Names command and analog, includes runnable example 2026.09.15

03 **Full page info in v2p, p2v, pageinfo** IMPROVED 60

Expands the `v2p`, `p2v`, and `pageinfo` commands to display full page info including vaddr, paddr, and page.

– Names commands and fields but no example given 2026.09.15

04 **Backwards disassembly in context and nearpc** NEW 60

Adds backwards disassembly in the context disasm view and `nearpc`, showing instructions before `$PC` when stopping at a breakpoint or after a syscall.

– Describes trigger conditions but lacks a usage example 2026.09.15

05 **pwndbg.dbg_mod.Error renamed to DebuggerError** BREAKING 60

`pwndbg.dbg_mod.Error` is renamed to `pwndbg.dbg_mod.DebuggerError`; any plugin or script that catches or references `pwndbg.dbg_mod.Error` will break on upgrade.

– Names exact symbol change and migration impact 2026.09.15

THINNER COVERAGE BELOW

06 **Improved indirect jump handling for x86** IMPROVED 55

Improves indirect jump handling for x86 PLT, `ret` instructions, and memory-indirect jumps that can segfault, enabling visualization of segfaulting ROP chains.

– Names affected instruction types, no direct command shown 2026.09.15

07 **Emulation for x64 kernel debugging** NEW 50

Enables emulation when debugging x64 kernels, unlocking emulated instruction flow and concrete operand display in the disassembly system.

– Explains benefit but no command or config surface named 2026.09.15

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! `pwndbg.dbg_mod.Error` is renamed to `pwndbg.dbg_mod.DebuggerError`; any plugin or script that catches or references `pwndbg.dbg_mod.Error` will break on upgrade.

> 2026.09.15 NOTES ↗ CODE ↗ LATEST STABLE
>
> pwndbg gains backwards disassembly, new `stack-vis` command, `nearpc -f` without symbols, and full page info on `v2p` / `p2v` / `pageinfo`

└──▷ GET THIS VERSION

```
$ git clone --branch 2026.09.15 https://github.com/pwndbg/pwndbg.git
# already have the repo? check out this version:
$ git checkout 2026.09.15
```

└──▷ TRY IT

Visualize the current thread's stack frames at a breakpoint to understand frame layout during exploit development.

```
$ stack-vis
```

Disassemble an entire function by address without debug symbols, getting branch visualization and operand annotations that raw `disass` lacks.

```
$ nearpc --function 0x401080
```

Adds `stack-vis` command to visualize stack frames of the current thread, similar to how `vis-heap-chunks` visualizes heap chunks.

Expands `nearpc -f` / `--function` to work without debugging symbols, making it a full replacement for GDB's/LLDB's `disass` command with branch visualization, branch targets, and operand annotations.

Expands `v2p`, `p2v`, and `pageinfo` commands to display full page info including vaddr, paddr, and page.

Adds backwards disassembly in context disasm and `nearpc`, showing instructions before `$PC` when stopping at a breakpoint or after a syscall.

Enables emulation when debugging x64 kernels, unlocking emulated instruction flow and concrete operand display in the disassembly system.

+1 more

└──▷ BREAKING ON UPGRADE

! `pwndbg.dbg_mod.Error` is renamed to `pwndbg.dbg_mod.DebuggerError`; any plugin or script that catches or references `pwndbg.dbg_mod.Error` will break on upgrade.

WAS THIS USEFUL? 👍 👎

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.9.1 NOTES ↗ CODE ↗ RANK

A framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap v0.9.1 adds an MCP server for LLM-driven PCAP analysis, structured Modbus and S7comm OT/ICS threat-hunting with CISA-mapped rules, and Hyperscan-accelerated pattern matching, alongside service-mode, database, and web UI improvements.

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

01 **MCP server for LLM-driven PCAP analysis** NEW 85

Adds a Model Context Protocol (MCP) server exposing tools including `query_audit_records`, `get_chart_data`, `extract_iocs_from_file`, `get_image_exif`, `lookup_cve`, `resolve`, `lookup_dhcp_fingerprint`, `get_live_capture_status`, `stop_live_capture`, `delete_session`, and `wait_for_session`. `lookup_cve` fetches real NVD vulnerability data with caching, and the server exposes a `netcap://session/{id}/graph` resource, a Prometheus metrics endpoint at `/mcp/metrics`, and `?sessionId` / `?inputFile` query parameters for per-request session selection shared with the web UI.

— Names every tool and endpoint but no runnable example given v0.9.1

02 **S7comm and S7CommPlus OT threat hunting** NEW 83

Adds OT/ICS threat-hunting rules for Siemens S7comm mapped to CISA advisory AA26-231A, covering write, logic download, logic theft, PLC stop, and CPU restart function codes, shipped as `rules/examples/s7comm_hunt.yml`. Also adds S7CommPlus cleartext framing parsing and a `PayloadObscured` flag to the S7comm stream decoder.

Hunt for unauthorized S7comm operations (logic download, PLC stop) on an OT network PCAP, with an allowlist of known-good engineering workstations to suppress false positives.

```
$ ./net capture -read ot-traffic.pcap --rules
rules/examples/s7comm_hunt.yml --approved-workstations
10.0.1.50,10.0.1.51
```

— Named rule file, CISA mapping, and a runnable hunt command v0.9.1

03 **Hyperscan/Vectorscan acceleration** NEW 78

Adds an optional Hyperscan/Vectorscan prefilter for nmap service probes (~2.2x faster) and CMS/web framework detection (~1.4x), plus a Hyperscan-backed rule-engine

`MatchesPattern` operator (up to ~6x faster on miss-heavy traffic); enabled by building with `-tags hyperscan`.

— Concrete speedups and build tag given, no full command v0.9.1

04 **Structured Modbus decoding and RTU support** NEW 75

Adds `--modbus-rtu-endpoints` (env `NC_MODBUS_RTU_ENDPOINTS`) to allowlist Modbus RTU-over-TCP destinations (IP:port or [IPv6]:port) for stream decoding, plus structured Modbus PDU decoding with threat-hunting rules covering all remaining Modbus function codes, pairing reads with writes via ordered correlation.

— Flag and env var named but no worked example shown v0.9.1

05 **Service mode preload and database rebuild controls** NEW 75

Adds `--service-preload-dir` (env `NC_SERVICE_PRELOAD_DIR`, default `<service-data-dir>/pcaps`) to preload a directory of PCAPs when running in service mode, `NC_DBS_SKIP_SOURCES` to skip specific data sources during database rebuild, and a `/health` readiness endpoint on the database (`dbs`) HTTP listener available before the initial rebuild completes.

Start service mode and preload a directory of historical PCAPs for immediate LLM-driven analysis via the MCP server.

```
$ ./net capture --service --service-preload-dir
/data/forensics/pcaps
```

— Flags, env vars and endpoint named with a runnable example v0.9.1

06 **Time and behavior helpers in detection rule engine** NEW 65

Adds source-to-distinct-destination cardinality (fan-out) detection and record-time thresholds to the detection rule engine, along with `IsBusinessHours` and `HourOfDay` time-of-day expression helpers and an `IsApprovedWorkstation` helper backed by a new `--approved-workstations` allowlist flag.

— Helpers and flag named but no dedicated example v0.9.1

THINNER COVERAGE BELOW

07 **Dashboard and fingerprints API in web UI** NEW 55

Adds a configurable dashboard landing page with built-in security views to the web UI, and server-side filter, sort, and pagination plus a per-row community-ID cap to the `/api/fingerprints` endpoint.

— Endpoint named but no request example or UI path v0.9.1

08 **Worker-owned TCP pools for concurrent capture** IMPROVED 45

Replaces the global assembly lock with per-worker TCP pool ownership, enabling concurrent capture.

— Describes the mechanism but gives no config or flag v0.9.1

09 **Deterministic offline capture output** IMPROVED 35

Offline capture output is now deterministic: the same PCAP produces identical audit records across runs.

— States the before/after but no mechanism detail v0.9.1

10 **GeoIP enrichment for Connection records** NEW 25

Adds GeoIP enrichment to Connection records.

— Single-line description with no mechanism or fields named v0.9.1

11 **Fuzz targets for packet and stream parsers** NEW 20

Adds 19 fuzz targets for packet and stream parsers.

— Bare count with no target names or usage v0.9.1

v0.9.1 NOTES ↗ CODE ↗ LATEST STABLE

Netcap v0.9.1 adds an MCP server for LLM-driven PCAP analysis, Modbus threat-hunting rules, S7comm OT detection, and Hyperscan acceleration.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.9.1 https://github.com/dreadl0ck/netcap.git
# already have the repo? check out this version:
$ git checkout v0.9.1
```

└──▷ TRY IT

Hunt for unauthorized S7comm operations (logic download, PLC stop) on an OT network PCAP, with an allowlist of known-good engineering workstations to suppress false positives.

```
$ ./net capture -read ot-traffic.pcap --rules
rules/examples/s7comm_hunt.yml --approved-workstations
10.0.1.50,10.0.1.51
```

Start service mode and preload a directory of historical PCAPs for immediate LLM-driven analysis via the MCP server.

```
$ ./net capture --service --service-preload-dir
/data/forensics/pcaps
```

Adds `--modbus-rtu-endpoints` flag (env: `NC_MODBUS_RTU_ENDPOINTS`) to specify a comma-separated allowlist of RTU-over-TCP destinations (IP:port or [IPv6]:port) for Modbus RTU stream decoding.

Adds `--service-preload-dir` flag (env: `NC_SERVICE_PRELOAD_DIR`) to specify a directory of PCAPs to preload when running in service mode (defaults to `<service-data-dir>/pcaps`).

Adds `--approved-workstations` flag for loading an approved-workstation allowlist used by detection rules via the `IsApprovedWorkstation` expression helper.

Adds an MCP (Model Context Protocol) server for LLM-driven PCAP analysis, exposing tools including `query_audit_records`, `get_chart_data`, `extract_iocs_from_file`, `get_image_exif`, `lookup_cve`, `resolve`, `lookup_dhcp_fingerprint`, `get_live_capture_status`, `stop_live_capture`, `delete_session`, and `wait_for_session`.

Adds `lookup_cve` MCP tool that fetches real NVD vulnerability data with caching.

+17 more

WAS THIS USEFUL? 👍 👎

NETRESEC

Feedback ⓘ Unmasking Malware Families (2026-09-21)

BLOG ↗ RANK

Network forensics, packet sniffers and IT security products. Download NetworkMiner and other free software for network security analysis.

> NetResec introduced FlowCarp, a statistical C2 protocol classifier that identifies malware families from packet behaviour rather than signatures, shipping both a CLI for training custom models and a hosted demo API.

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

01 FlowCarp CLI for statistical C2 protocol classification NEW

95

FlowCarp is a new statistical application-layer protocol classifier that identifies C2 protocols — including encrypted and proprietary ones — by analyzing packet sizes, timing, message length distributions, and request/response behavior rather than byte-sequence signatures. The `flowcarp` CLI adds `--input`, `--output`, and `--model` flags for building and applying custom protocol models from PCAP files: `flowcarp --input <pcap> --output model:<label>:<class>:<file>.fcp` trains a new model from example C2 capture sessions, and `flowcarp --model <file>.fcp --input <pcap> --output -` classifies traffic in a new PCAP against a previously built custom model.

Train a custom FlowCarp model from sandbox-captured C2 sessions so you can classify unknown malware traffic without writing IDS signatures.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Classify traffic from a new sandbox capture against your custom protocol model to confirm C2 family attribution.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

— Names exact flags, commands and file format for training/classifying

Unmasking Malware Families (2026-09-21)

02 Hosted FlowCarp demo REST API NEW

70

A demo REST endpoint at `https://demo.flowcarp.com` accepts raw PCAP via HTTP POST and returns a JSON response of detected flows, protocols, and alerts. The hosted

demo service detects more than 600 protocols and supports adding custom protocol models built from user-supplied example traffic.

Submit a PCAP to the FlowCarp demo service and pipe the JSON results into jq for quick triage during incident response.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

— Endpoint and response format named but limited mechanism detail

Unmasking Malware Families (2026-09-21)

Unmasking Malware Families (2026-09-21) BLOG ↗ LATEST STABLE

FlowCarp brings statistical C2 protocol classification to PCAP analysis, identifying 600+ malware families without signature rules.

└──▷ **TRY IT**

Train a custom FlowCarp model from sandbox-captured C2 sessions so you can classify unknown malware traffic without writing IDS signatures.

```
$ flowcarp --input xyz.pcap --output
model:MalwareXYZ:MALICIOUS:malwarexyz.fcp
```

Classify traffic from a new sandbox capture against your custom protocol model to confirm C2 family attribution.

```
$ flowcarp --model malwarexyz.fcp --input new.pcap --output -
```

Submit a PCAP to the FlowCarp demo service and pipe the JSON results into jq for quick triage during incident response.

```
$ curl --data-binary @dump.pcap https://demo.flowcarp.com | jq
```

Adds `--input`, `--output`, and `--model` flags to the `flowcarp` CLI for building and applying custom statistical protocol models from PCAP files.

Adds `flowcarp --input <pcap> --output model:<label>:<class>:<file>.fcp` to train a new protocol model from example C2 capture sessions.

Adds `flowcarp --model <file>.fcp --input <pcap> --output -` to classify traffic in a new PCAP against a previously built custom model.

> Exposes a demo REST endpoint at `https://demo.flowcarp.com` accepting raw PCAP via HTTP POST and returning a JSON response of detected flows, protocols, and alerts.
>
> Introduces FlowCarp, a statistical application-layer protocol classifier that identifies C2 protocols — including encrypted and proprietary ones — by analyzing packet sizes, timing, message length distributions, and request/response behavior rather than byte-sequence signatures.
>
> +1 more

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ announced · Sep 21, 2026 PRESS ↗ RANK

Fastly's edge cloud platform delivers faster, safer, and more scalable sites and apps to customers. Elevate your edge CDN, video delivery, security, and more.

> Fastly launched an AI-focused security and governance stack at the edge: AI Runtime Control (ARC) as a managed control plane for LLM traffic, an AI Firewall to catch prompt injection attacks, and API Security enforcement extended to agentic AI traffic.

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

01 AI Runtime Control (ARC) for LLM traffic NEW

77

ARC is a Fastly-managed control plane that routes every model call through a single endpoint across public and self-hosted AI providers, providing access governance, real-time token spend visibility, rate limiting, and resilience. Per-app and per-developer virtual keys can be created, rotated, or revoked instantly without client-side code changes, replacing shared provider API keys and protecting underlying provider credentials; enforceable monthly dollar-budget limits can be applied directly to individual virtual keys. Automatic failover lets a virtual key point at an ordered list of model targets (primary, fallback, and further) so a provider outage is transparent to callers, and ARC delivers full request and response logging for every AI interaction (prompts and completions), searchable by user, model, provider, and date range. It operates as bring-your-own-key (BYOK), preserving existing provider relationships and negotiated rates.

– Rich mechanism and named controls but no API/CLI syntax given launch-20260921-de13b213 launch-20260921-0c070049

THINNER COVERAGE BELOW

02 AI Firewall for prompt injection detection NEW

55

AI Firewall evaluates prompts directly in the request path at the edge to detect and block LLM-based attacks, including prompt injection, before they reach targeted models. It is powered by an updated version of Fastly's SmartParse detection engine and, within ARC, supports optional custom guardrails.

– Names detection engine but no config or endpoint given launch-20260921-de13b213 launch-20260921-0c070049

03 API Security enforcement for agentic AI traffic NEW

45

Fastly's API Security now applies API contract rules to agentic, agent-assisted, and conventional traffic on a service-by-service basis, letting organizations observe or block non-conforming requests from AI agents accessing enterprise APIs.

— Describes scope but no concrete config or rule syntax launch-20260921-0c070049

announced · Sep 21, 2026 PRESS ↗

Fastly launches AI Runtime Control, AI Firewall, and expanded API Security to govern, protect, and control AI traffic at the edge.

Adds AI Runtime Control, which routes model calls through a single endpoint across public and self-hosted AI providers, with virtual keys to protect underlying provider credentials, real-time token spend visibility, rate limiting, budget controls, and failover.

Adds AI Firewall, which evaluates prompts directly in the request path at the edge to detect and block LLM-based attacks including prompt injection before they reach targeted models.

Adds API Security enforcement for agentic traffic, applying API contract rules to agentic, agent-assisted, and conventional traffic service by service — allowing organizations to observe or block non-conforming requests from AI agents accessing enterprise APIs.

announced · Sep 21, 2026 BLOG ↗

Fastly launches AI Runtime Control (ARC), a control plane for governing, securing, and auditing all LLM provider traffic.

Introduces AI Runtime Control (ARC), a Fastly-managed control plane that routes every LLM model call through a single endpoint for access governance, spend attribution, and resilience.

Provides per-app and per-developer virtual keys that can be created, rotated, or revoked instantly without code changes on the client side, replacing shared provider API keys.

Applies enforceable monthly dollar-budget limits directly to individual virtual keys for proactive AI spend control.

Supports automatic failover by letting a virtual key point at an ordered list of model targets (primary, fallback, and further) so a provider outage is transparent to callers.

Includes AI Firewall within ARC — powered by an updated version of Fastly's SmartParse detection engine — to detect prompt injection and other LLM-specific attacks via optional custom guardrails.

+2 more

WAS THIS USEFUL?  

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Falco

Feedback ⓘ 0.45.0 NOTES ↗ RANK

Cloud Native Runtime Security

> Falco 0.45 reworks rule matching to operate on raw bytes and adds new multi-value string modifiers, alongside a reload-status API endpoint and several smaller CLI and Helm chart additions.

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

01 GET /reload API and Unix socket reload control NEW

90

Exposes reload status via `GET /reload` and adds optional reload control over a protected Unix socket on Linux, preserving SIGHUP reception and pending requests across reloads. Status can be checked with `curl -s --unix-socket /run/falco/falco.sock http://localhost/reload`.

Check whether Falco has finished reloading after sending SIGHUP, without restarting the process.

```
$ curl -s --unix-socket /run/falco/falco.sock
http://localhost/reload
```

— Runnable curl command and exact endpoint path given 0.45.0

02 Raw-byte rule matching and regex sanitization BREAKING

85

Rule conditions now evaluate on raw field bytes, enabling matching of invalid and non-printable UTF-8 sequences with string operators and single-byte matching via the new `\xHH` escape sequence (e.g. `evt.buffer contains "\x00"`). The `regex` operator remains an exception: it operates on sanitized input, replacing invalid UTF-8 bytes with `U+FFFD` before matching, while non-printable characters are escaped at output encoding time. This is a breaking change: conditions using non-regex operators to match `U+FFFD` in place of invalid bytes will no longer behave the same and must be updated to match the original raw bytes; `FALCO_ENGINE_VERSION` is bumped from `0.62.0` to `0.63.0`.

Catch filenames containing a raw null byte or other non-printable byte (e.g. evasion via NUL-injection) using the new `\xHH` hex escape in a byte-buffer condition.

```text
evt.buffer contains "\x00"
```

— Names exact escape syntax and version bump, but no migration script `0.45.0`

---

03 **anyof/allof/oneof string comparison modifiers** NEW

75

Adds `anyof`, `allof`, and `oneof` string comparison modifiers (since Falco 0.44) so a single condition can test one field against multiple values or patterns, e.g. proc.cmdline contains anyof (curl, wget) or proc.name != allof (bash, sh). Also adds support for referencing a Falco list inside the modifier parentheses, e.g. fd.name contains anyof (suspicious_paths).

Detect processes invoking common exfiltration tools without writing a separate rule for each binary, using the `anyof` modifier to match multiple values in one condition.

```text
proc.cmdline contains anyof (curl, wget, nc, scp)
```

Exclude multiple known-good shell names in a single condition using `allof` with `!=`, rather than chaining separate inequality checks.

```env
proc.name != allof (bash, sh)
```

— Concrete syntax with runnable rule conditions given `product docs`

---

04 **container_engines config key deprecation warning** DEPRECATED

60

Configurations still setting the removed `container_engines` key will now log a schema validation warning at startup, since the key has been dropped from the configuration JSON schema.

— Names the exact key removed but no replacement guidance `0.45.0`

THINNER COVERAGE BELOW

05 **serviceAccount.labels in Helm chart** NEW

50

Adds `serviceAccount.labels` to the Helm chart to set custom labels on the ServiceAccount.

— Names the config key but no further mechanism or example `0.45.0`

---

06 **--dry-run suggestion on --validate** IMPROVED 40

Falco now suggests `--dry-run` when `--validate` receives a Falco configuration file.

— Names both flags but no example output shown `0.45.0`

└──▷ **BREAKING ON UPGRADE**

! When upgrading from Falco 0.44 or earlier, conditions that used non-regex operators to match the Unicode replacement character `U+FFFD` in place of invalid bytes will no longer behave the same — those conditions must be updated to match the original raw bytes instead.

! Rules that matched the UTF-8 replacement character using non-regex operators must be reviewed: rule conditions now operate on raw field bytes, and `FALCO_ENGINE_VERSION` is bumped from `0.62.0` to `0.63.0`.

! Configurations still setting the removed `container_engines` key will now log a schema validation warning at startup (the key has been dropped from the configuration JSON schema).

`0.45.0` NOTES ↗ LATEST STABLE

Falco 0.45.0 adds a `GET /reload` API endpoint, raw-byte rule matching with `\xHH` escape sequences, and a `--dry-run` suggestion on `--validate`.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 0.45.0
https://github.com/falcosecurity/falco.git
# already have the repo? check out this version:
$ git checkout 0.45.0
```

└──▷ **TRY IT**

Check whether Falco has finished reloading after sending SIGHUP, without restarting the process.

```
$ curl -s --unix-socket /run/falco/falco.sock
http://localhost/reload
```

Exposes reload status via `GET /reload` and adds optional reload control over a protected Unix socket on Linux, preserving SIGHUP reception and pending requests across reloads.

Evaluates rule conditions on raw field bytes, enabling matching of invalid and non-printable UTF-8 sequences with string operators and single-byte matching with the new `\xHH` escape sequence; `regex` now operates on sanitized input and non-printable characters are escaped at output encoding time.

Adds `serviceAccount.labels` to the Helm chart to set custom labels on the ServiceAccount.

Suggests `--dry-run` when `--validate` receives a Falco configuration file.

└──▷ **BREAKING ON UPGRADE**

- **!** Rules that matched the UTF-8 replacement character using non-regex operators must be reviewed: rule conditions now operate on raw field bytes, and `FALCO_ENGINE_VERSION` is bumped from `0.62.0` to `0.63.0`.
- **!** Configurations still setting the removed `container_engines` key will now log a schema validation warning at startup (the key has been dropped from the configuration JSON schema).

WAS THIS USEFUL? 👍 👎

Alibaba Cloud OpenSandbox

Feedback ⓘ release-1.1.0

NOTES ↗ RANK

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

> OpenSandbox 1.1.0 ships Firecracker FastPath fleets, an opt-in PostgreSQL snapshot store with HA coordination, sandbox lifecycle hooks, and a scoped diagnostics API, alongside a wide set of egress, security, and Kubernetes provider hardening changes — plus breaking migrations to `agents.x-k8s.io/v1beta1` and removal of the `informer_enabled` flag.

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

01 **Scoped diagnostics API** NEW 90

Adds a stable, scoped diagnostics API with `logs:container|all` and `events:runtime|lifecycle|all` scopes, returning `400 DIAGNOSTICS_SCOPE_UNSUPPORTED` for unknown scopes and backward-compatible plain-text responses.

Scope diagnostics queries to only runtime events for a sandbox, using the new `events:runtime` scope to cut noise during incident triage.

```
$ curl -H 'OPEN-SANDBOX-API-KEY: <key>'
'https://<server>/diagnostics?scope=events:runtime'
```

— Runnable curl example with exact scopes and error code release-1.1.0

02 **PostgreSQL snapshot store with HA coordination** NEW 88

Adds an opt-in PostgreSQL snapshot store with pooled connections and CAS semantics, a one-shot `migrate-snapshots` command for moving existing SQLite records, and multi-process HA coordination for snapshot create/recovery/delete on Kubernetes.

Enable the PostgreSQL snapshot store for HA multi-process deployments and migrate existing SQLite records in one shot.

```
$ migrate-snapshots
```

— Runnable migrate-snapshots command included release-1.1.0

03 **Egress sidecar resource, OTLP export, and timeout config** NEW

85

Adds `[egress.resources]` for independent resource requests/limits on the Kubernetes egress sidecar, `[egress] otlp_endpoint` so the sidecar exports metrics to an OpenTelemetry collector on both runtimes, and `[egress] timeout_seconds` to make the sidecar's readiness wait configurable.

Export egress sidecar metrics to your OpenTelemetry collector and cap its resource footprint independently of the sandbox container.

```toml
[egress]
otlp_endpoint = "http://otel-collector:4317"
timeout_seconds = 30

[egress.resources]
requests = { cpu = "50m", memory = "64Mi" }
limits   = { cpu = "200m", memory = "128Mi" }
```

— Exact config keys with runnable TOML example `release-1.1.0`

---

04 **Sandbox lifecycle hooks (OSEP-0020)** NEW

80

Adds `preStart` and `periodic` lifecycle hooks to sandbox create requests, validated server-side and serialized into a reserved environment variable for Kubernetes workloads, with `preStart.timeoutSeconds` raised to 10800 seconds. Lifecycle hook models were added to the API spec and every SDK with validation consistent with server-side ranges.

— Names field, env var, and timeout limit but no example call `release-1.1.0`

---

05 **Pool-only sandbox creation and back-pressure signaling**

IMPROVED

80

Supports pool-only sandbox creation using `poolRef` without an entrypoint (previously rejected with HTTP 400), and surfaces pool back-pressure as HTTP 429 `KUBERNETES::POOL_CAPACITY_EXHAUSTED` with a `Retry-After` header.

— Exact status codes and field named, no curl example `release-1.1.0`

---

06 **agent-sandbox provider migrates to v1beta1** BREAKING

75

The agent-sandbox provider and ingress informer migrate to `agents.x-k8s.io/v1beta1` (required by agent-sandbox v1.0.0), creating workloads with `operatingMode: Running` and `service: true`; existing workloads using the prior API version will break.

— Exact API version and fields named, no migration steps `release-1.1.0`

07 **Docker runtime config parity with Kubernetes** NEW 70

Adds `sandbox_env` and `sandbox_binds` Docker-runtime-wide configuration keys, giving Docker parity with the Kubernetes pod template for fleet-wide CA bundles, proxies, and mounts.

— Named config keys, no example command release-1.1.0

08 **FastPath fleets lifecycle backend** NEW 70

Adds the FastPath fleets lifecycle backend: create, renew, delete, and status mapping via FastPath v2, plus end-to-end fast-sandbox pause/resume with async 202 semantics and UID fencing.

— Mechanism named, no exact command shown release-1.1.0

09 **Secure-access signing keys from env/Secret** IMPROVED 65

Sources OSEP-0011 secure-access signing keys from the `OPENSANDBOX_SECURE_ACCESS_KEYS` environment variable or a Kubernetes Secret instead of plaintext chart values.

— Named env var, no walkthrough of migration release-1.1.0

10 **Per-sandbox image pull secrets** NEW 65

Honors `image.auth` in create requests for the agent-sandbox provider, with per-sandbox pull secrets garbage-collected along with the Sandbox CR.

— Named field and GC behavior, no example release-1.1.0

11 **Configurable proxy internal resolution** NEW 60

Adds `[proxy] resolve_internal` config key to make the reverse proxy's internal resolution configurable for deployments that cannot route to Docker bridge IPs.

— Named config key but no example given release-1.1.0

12 **gVisor snapshot rejection** NEW 60

Rejects unsupported gVisor (`runsc`) snapshots before persistence with `409 SNAPSHOT::UNSUPPORTED_RUNTIME` .

— Exact error code named, no reproduction example release-1.1.0

THINNER COVERAGE BELOW

13 **Server-side HTTP request metrics over OTLP** NEW 55

Exports generic HTTP request metrics (request rate, error rate, latency) from the server over OTLP, complementing the SDK-reported create-duration histogram.

— Named metrics and OTLP but no endpoint/config key release-1.1.0

14 **Strict CodeInterpreter health check by default** BREAKING

55

Defaults `CodeInterpreter` creation to a strict health check in every SDK language, blocking until both execd and the Jupyter kernel gateway are verified ready.

— Explains new default behavior, no opt-out flag given release-1.1.0

15 **Egress telemetry spoofing prevention** IMPROVED

50

Injects `OPENSANDBOX_EGRESS_SANDBOX_ID` server-side so egress sidecar telemetry can no longer be spoofed by request environment variables.

— Explains fix but no reproduction steps release-1.1.0

16 **informer_enabled flag removed** BREAKING

50

The `informer_enabled` flag is removed; the informer-backed cache is now the only Kubernetes read path and the flag can no longer be used to disable it.

— Named flag and removal, no replacement guidance release-1.1.0

17 **FastPath snapshot and template management in SDKs** NEW

40

Adds fast-sandbox snapshot support and fsb template management across all five SDKs.

— Thin description, no SDK method names given release-1.1.0

└──▷ BREAKING ON UPGRADE

! The agent-sandbox provider and ingress informer migrate to `agents.x-k8s.io/v1beta1` (required by agent-sandbox v1.0.0), creating workloads with `operatingMode: Running` and `service: true`; existing workloads using the prior API version will break.

! The `informer_enabled` flag is removed; the informer-backed cache is now the only Kubernetes read path and the flag can no longer be used to disable it.

release-1.1.0 NOTES ↗ LATEST STABLE

OpenSandbox 1.1.0 debuts Firecracker FastPath sandboxes, a PostgreSQL snapshot store, lifecycle hooks, and a scoped diagnostics API

└──▷ GET THIS VERSION

```
$ git clone --branch release-1.1.0 https://github.com/opensandbox-group/OpenSandbox.git
# already have the repo? check out this version:
$ git checkout release-1.1.0
```

└──▷ **TRY IT**

Enable the PostgreSQL snapshot store for HA multi-process deployments and migrate existing SQLite records in one shot.

```
$ migrate-snapshots
```

Scope diagnostics queries to only runtime events for a sandbox, using the new `events:runtime` scope to cut noise during incident triage.

```
$ curl -H 'OPEN-SANDBOX-API-KEY: <key>'
'https://<server>/diagnostics?scope=events:runtime'
```

Export egress sidecar metrics to your OpenTelemetry collector and cap its resource footprint independently of the sandbox container.

```toml
[egress]
otlp_endpoint = "http://otel-collector:4317"
timeout_seconds = 30

[egress.resources]
requests = { cpu = "50m", memory = "64Mi" }
limits   = { cpu = "200m", memory = "128Mi" }
```

Adds `preStart` and `periodic` lifecycle hooks to sandbox create requests (OSEP-0020), validated server-side and serialized into a reserved environment variable for Kubernetes workloads; `preStart.timeoutSeconds` limit raised to 10800 seconds.

Adds opt-in PostgreSQL snapshot store with pooled connections and CAS semantics, plus the one-shot `migrate-snapshots` command for moving existing SQLite records, and multi-process HA coordination for snapshot create/recovery/delete on Kubernetes.

Adds `[egress.resources]` config key for independent resource requests/limits on the Kubernetes egress sidecar.

Adds `[egress] otlp_endpoint` config key so the egress sidecar exports metrics to an OpenTelemetry collector on both runtimes.

Adds `[egress] timeout_seconds` config key to make egress sidecar readiness wait configurable.

+13 more

└──▷ **BREAKING ON UPGRADE**

! The agent-sandbox provider and ingress informer migrate to `agents.x-k8s.io/v1beta1` (required by agent-sandbox v1.0.0), creating workloads with `operatingMode: Running` and `service: true`; existing workloads using the prior API version will break.

! The `informer_enabled` flag is removed; the informer-backed cache is now the only Kubernetes read path and the flag can no longer be used to disable it.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ API keys now hold a role: NOTES ↗ RANK

Plerion is a cloud security platform that discovers and manages vulnerabilities and misconfigurations across infrastructure.

> Plerion introduced role-based access control for tenant API keys, including a role-scoped OpenAPI discovery endpoint and a new permission gate on IaC scan uploads, and shipped five new or expanded AWS security checks covering Redshift VPC routing, Security Hub AI best practices, Kinesis encryption, Organizations exit controls, and ALB access logging.

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

01 **Permission gate on IaC scan uploads** BREAKING 88

IaC scan uploads are now gated behind a new `CodeRepository:RequestScan` permission on `POST /v1/tenant/shiftleft/iac/scan`; custom roles can add this permission via Request code scans under Manage integrations.

Upload an IaC scan result using a key whose role carries the `CodeRepository:RequestScan` permission.

```
$ curl -s -X POST
https://api.plerion.com/v1/tenant/shiftleft/iac/scan \
  -H 'Authorization: Bearer <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d @scan-results.json
```

— Named endpoint, permission and UI path with curl example API keys now hold a role:

02 **Role-scoped OpenAPI reference endpoint** NEW 81

New `GET /v1/tenant/openapi` endpoint returns the API reference scoped to the calling key's role, letting a key discover exactly what it is permitted to call.

Audit what API operations a specific key is allowed to call — useful when scoping a least-privilege integration key.

```
$ curl -s -H 'Authorization: Bearer <your-api-key>'
https://api.plerion.com/v1/tenant/openapi
```

— Named endpoint with a runnable curl example API keys now hold a role:

03 **Redshift enhanced VPC routing check** NEW 80

New check `PLERION-AWS-329` (High severity) fails any Redshift cluster whose `EnhancedVpcRouting` setting is off, ensuring COPY and UNLOAD traffic to S3 stays inside the VPC and is visible to security groups, network ACLs, endpoint policies, and VPC Flow Logs. Maps to AWS Foundational Security Best Practices Redshift.7, CIS Controls v8, and NIST 800-53.

— Named check ID and framework mappings, no remediation command

Say hello to a new Redshift check:

04 **Security Hub AI Best Practices standard check** NEW 80

New check `PLERION-AWS-1177` (Medium severity) fails any AWS region where the Security Hub AI Security Best Practices standard is not in `READY` state, covering 31 automated controls for Amazon Bedrock, Amazon Bedrock AgentCore, and Amazon SageMaker. Maps to the AWS Well-Architected Machine Learning Lens and Generative AI Lens.

— Named check ID and scope, no remediation steps given Say hello to a new Redshift check:

05 **Kinesis Data Streams encryption check** NEW 73

New check `PLERION-AWS-1176` fails any Kinesis Data Streams channel encrypted with an AWS-owned key rather than a customer-managed KMS key; Plerion now collects Kinesis channels as a new asset type.

— Named check ID and new asset type, no remediation command

Say hello to a new Redshift check:

06 **Organization member-account exit control check** NEW 73

New check `PLERION-AWS-1178` (High severity) fails any AWS organization where at least one active member account is not covered by an unconditional SCP Deny of both `organizations:LeaveOrganization` and `account:CloseAccount`.

— Named check ID and exact SCP actions required Say hello to a new Redshift check:

07 **Role assignment for API keys** NEW 65

Tenant API keys now carry a role: new keys can be created with Tenant read-only or a custom role, and an existing key's role can be changed from the API keys list in the console.

Assign a least-privilege role to an API key at creation time so it cannot perform write operations.

📍 In the Plerion console, go to API keys › Create key, set Role to Tenant read-only, then save. To update an existing key, select it from the API keys list and

choose a different role.

— UI steps given but no API endpoint for the update itself `API keys now hold a role:`

---

08 **ALB access logging check now recognizes CloudWatch Logs**

IMPROVED 62

Check `PLERION-AWS-791` ('Application Load Balancer does not have access logging enabled') now passes ALBs that deliver access logs to CloudWatch Logs in addition to those logging to S3.

— Named check ID and before/after but no remediation steps

`Say hello to a new Redshift check:`

`API keys now hold a role:` NOTES ↗ STABLE

Plerion API keys now carry a role, and a new IaC scan upload permission gates POST `/v1` /tenant/shiftleft/iac/scan.

└──▷ **TRY IT**

Audit what API operations a specific key is allowed to call — useful when scoping a least-privilege integration key.

```
$ curl -s -H 'Authorization: Bearer <your-api-key>'
https://api.plerion.com/v1/tenant/openapi
```

Upload an IaC scan result using a key whose role carries the `CodeRepository:RequestScan` permission.

```
$ curl -s -X POST
https://api.plerion.com/v1/tenant/shiftleft/iac/scan \
  -H 'Authorization: Bearer <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d @scan-results.json
```

Assign a least-privilege role to an API key at creation time so it cannot perform write operations.

📍 In the Plerion console, go to API keys › Create key, set Role to `Tenant read-only`, then save. To update an existing key, select it from the API keys list and choose a different role.

Adds role assignment to every tenant API key; new keys can be created with `Tenant read-only` or a custom role, and an existing key's role can be changed from the API

keys list.

Adds `GET /v1/tenant/openapi` endpoint that returns the API reference scoped to the calling key's role, so a key can discover exactly what it is permitted to call.

Gates IaC scan uploads behind a new `CodeRepository:RequestScan` permission on `POST /v1/tenant/shiftleft/iac/scan`; custom roles can add it via `Request code scans` under `Manage integrations`.

Say hello to a new Redshift check: NOTES ↗ STABLE

Plerion adds four new checks: Redshift enhanced VPC routing, Security Hub AI Best Practices, Kinesis channel encryption, and Organizations member-account exit controls.

Adds check `PLERION-AWS-329` (High severity) that fails any Redshift cluster whose `EnhancedVpcRouting` setting is off, ensuring COPY and UNLOAD traffic to S3 stays inside the VPC and is visible to security groups, network ACLs, endpoint policies, and VPC Flow Logs; maps to AWS Foundational Security Best Practices Redshift.7, CIS Controls v8, and NIST 800-53.

Adds check `PLERION-AWS-1177` (Medium severity) that fails any AWS region where the Security Hub AI Security Best Practices standard is not in `READY` state, covering 31 automated controls for Amazon Bedrock, Amazon Bedrock AgentCore, and Amazon SageMaker; maps to the AWS Well-Architected Machine Learning Lens and Generative AI Lens.

Adds check `PLERION-AWS-1176` that fails any Kinesis Data Streams channel encrypted with an AWS-owned key rather than a customer-managed KMS key; Plerion now collects Kinesis channels as a new asset type.

Adds check `PLERION-AWS-1178` (High severity) that fails any AWS organization where at least one active member account is not covered by an unconditional SCP Deny of both `organizations:LeaveOrganization` and `account:CloseAccount`.

Expands check `PLERION-AWS-791` ('Application Load Balancer does not have access logging enabled') to pass ALBs that deliver access logs to CloudWatch Logs in addition to those logging to S3.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Mobile

Feedback ⓘ 2026.38Released on September 21, 2026

NOTES ↗ RANK

Sophos Mobile is an enterprise mobile device management and security solution that protects smartphones and tablets through threat detection, compliance, and remote management.

> Sophos Mobile 2026.38 adds automatic MTD policy hardening for API-created policies, new Apple Business profile setup-assistant skip options, and support for the latest Apple OS versions.

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

01 New Apple Business profile setup-assistant settings NEW

60

Apple Business profiles gain three new settings to skip corresponding setup-assistant steps: Accessibility Appearance, Safety & Handling, and Liquid Glass. The Accessibility Appearance and Safety & Handling settings are iPhone/iPad only.

— Names exact settings but no walkthrough of use 2026.38Released on September 21, 2026

THINNER COVERAGE BELOW

02 Auto-config for Mobile API-created MTD policies IMPROVED

53

Sophos Mobile now automatically adds recommended configurations to Mobile Threat Defense policies created via the Mobile API. This applies only to newly created policies; existing policies are left unchanged.

— Explains scope and before/after but no config keys given

2026.38Released on September 21, 2026

03 iOS 27, iPadOS 27, and macOS 27 support NEW

40

Sophos Mobile adds support for iOS 27, iPadOS 27, and macOS 27 as managed device platforms.

— Bare platform support announcement with no detail on new capabilities

2026.38Released on September 21, 2026

2026.38Released on September 21, 2026 NOTES ↗ LATEST STABLE

> Sophos Mobile 2026.38 adds iOS/iPadOS/macOS 27 support, new Apple Business profile settings, and auto-config for Mobile API-created MTD policies.

> Automatically adds recommended configurations to Mobile Threat Defense policies created via the Mobile API (applies to new policies only; existing policies are unchanged).
>
> Adds `Accessibility Appearance`, `Safety & Handling`, and `Liquid Glass` settings to Apple Business profiles to skip the corresponding setup-assistant steps (Accessibility Appearance and Safety & Handling are iPhone/iPad only).
>
> Adds support for iOS 27, iPadOS 27, and macOS 27 as managed device platforms.

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Local response filters now officially released (2026-09-21) BLOG ↗ RANK

Trend Vision One is Trend Micro's cloud-based XDR platform providing unified security visibility and threat detection across endpoints, networks, and cloud environments.

> Trend Vision One reached general availability for local response filters that auto-terminate processes on Windows endpoints, and the Zero Trust Secure Access On-Premises Gateway added content inspection for Vertex Chat traffic.

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

01 **Local response filters for Windows endpoints** NEW 80

Local response filters are now generally available, automatically terminating processes on Windows endpoints upon detection to reduce mean time to detect. They can be managed via Endpoint Security → Endpoint Security Configuration → Endpoint Security Policies, activated per-detection from Agentic SIEM and XDR → Observed Attack Techniques, and all enabled filters can be listed under Agentic SIEM and XDR → Detection Model Management → Local response filters.

Audit which local response filters are currently active across your environment before enabling new ones.

> 📍 In the console, go to Agentic SIEM and XDR → Detection Model Management → Local response filters to view all enabled filters.

Embed automated process-termination responses into a Windows endpoint policy so detections are actioned without manual intervention.

> 📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Endpoint Security Policies, select the target policy, and enable the desired local response filters.

Trigger automated process termination directly from an observed attack technique event to shorten response time for a specific detection.

> 📍 In the console, go to Agentic SIEM and XDR → Observed Attack Techniques, locate the relevant technique event, and enable the associated local response filter.

— Exact console navigation paths given for each workflow

Local response filters now officially r…

THINNER COVERAGE BELOW

02 Vertex Chat content inspection in ZTSA gateway NEW

35

The Zero Trust Secure Access On-Premises Gateway (version 1.0.94) adds content inspection support for Vertex Chat traffic passing through the gateway.

— Names the traffic type but no mechanism or steps

Zero Trust Secure Access On-Premises Ga…

Local response filters now officially released (2026-09-21) BLOG ↗ STABLE

Trend Vision One GA: local response filters auto-terminate processes on Windows endpoints to cut MTTD.

└──▷ **HOW TO FIND IT**

Audit which local response filters are currently active across your environment before enabling new ones.

📍 In the console, go to Agentic SIEM and XDR → Detection Model Management → Local response filters to view all enabled filters.

Embed automated process-termination responses into a Windows endpoint policy so detections are actioned without manual intervention.

📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Endpoint Security Policies, select the target policy, and enable the desired local response filters.

Trigger automated process termination directly from an observed attack technique event to shorten response time for a specific detection.

📍 In the console, go to Agentic SIEM and XDR → Observed Attack Techniques, locate the relevant technique event, and enable the associated local response filter.

Enables local response filters for Windows endpoints that automatically terminate processes on detection, reducing mean time to detect.

Manages local response filters via Endpoint Security → Endpoint Security Configuration → Endpoint Security Policies to integrate detection into security policy workflows.

Activates local response filters per-detection in Agentic SIEM and XDR → Observed Attack Techniques to convert individual event detections into automated responses.

Lists all enabled local response filters in Agentic SIEM and XDR → Detection Model Management → Local response filters.

Zero Trust Secure Access On-Premises Gateway version 1.0.94 (2026-09-21)

BLOG ↗ LATEST STABLE

Zero Trust Secure Access On-Premises Gateway 1.0.94 adds Vertex Chat content inspection support.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

JumpCloud

Feedback ⓘ Release Notes 2022 (2026-09-21) BLOG ↗

RANK

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

> Across releases spanning 2022–2026, JumpCloud shipped its broadest identity and device-management expansion yet: Password Vault across mobile, Admin Portal, User Portal and browser extension; Password Manager and Remote Assist maturing into full-featured tools; JumpCloud Go passwordless auth, MCP Server, AI Commands Builder and Agent Identities; GA launches of Access Risk Detection and SaaS License Management; and breaking changes to the systemusers reactivate endpoint, Dynamic Groups filter operators, and LDAP TLS cipher suites.

└──▷ **WHAT SHIPPED · 64 FEATURES** most completely described first ? what's the number?

01 **Systemusers reactivate endpoint replaced** BREAKING 90

Replaces the deprecated `GET api/systemusers/reactivate/{id}` endpoint, which now returns a 404, with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` accepting `activation_email` in the request body; all callers must switch to the new endpoint.

Reactivate a user account after migrating scripts from the deprecated GET endpoint to the new POST endpoint.

```
$ POST https://console.jumpcloud.com/systemusers/reactivate/{id}
{
  "activation_email": "{email_address}"
}
```

— Exact old and new endpoints plus a runnable request body example

Release Notes 2024 (2026-09-21)

02 **Agent Identities for non-human AI agents** NEW 90

Introduces Agent Identities (Public Preview) under `Identity Management > Agents` to register and govern non-human AI agent accounts with OAuth-based authentication, group-based application access, and full lifecycle management independent of human user accounts. Adds a `client authentication type` setting on an agent's `Authentication` tab supporting `Client secret basic` (credentials in the

authorization header, default) and `Client secret post` (credentials in the request body) for OAuth token requests.

Register a new AI agent and scope its OAuth credentials to require credentials in the request body — needed when integrating with platforms like Amazon Bedrock AgentCore that enforce `client_secret_post`.

> 📍 In the Admin Portal: Go to Identity Management › Agents › + Add. Fill in agent name, owner, and company email on the Details tab. On the Authentication tab, generate OAuth client credentials, then set the client authentication type to `Client secret post` to send the client ID and secret in the request body.

— Named UI path, both OAuth modes, and a full setup example

Release Notes 2026 (2026-09-21)

---

03 **AI Gateway MCP visibility and dashboard** NEW 85

Adds an `MCP Connection` column (showing `Discovered`, `Connected`, or `Not Available`) in `Access > AI & SaaS Management > Applications` to surface MCP server status and add servers to AI Gateway without switching tools, plus a `Dashboard` tab on AI Gateway with org-wide MCP server usage metrics, tool call volume, error rates, trend charts, and per-server drill-down.

Surface which SaaS apps have locally running MCP servers and connect them to AI Gateway for centralized governance.

> 📍 In the Admin Portal: Go to Access › AI & SaaS Management › Applications. Review the MCP Connection column for apps showing `Discovered`. Click the status for a target app, review discovered MCP servers, then click `Add MCP Server` to register them in AI Gateway.

— Exact column values and navigation path with a worked example

Release Notes 2026 (2026-09-21)

---

04 **Password Vault across Admin Portal, User Portal, mobile app, and browser extension**

NEW 83

Launches Password Vault as a password-management product spanning four surfaces: the JumpCloud Protect mobile app (view secrets, autofill into mobile apps and mobile browser using Push/TOTP-authenticated sessions), the Admin Portal (enable the vault, assign access to users and user groups, manage shared credentials and folders with Connect/View/Edit/Manage permissions, and view audit logs in Directory Insights), the User Portal (manage and share credentials and website applications, bulk import via CSV or Excel, and organize items in personal and shared folders), and a Browser Extension

via JumpCloud Go (autofill of non-SSO website passwords and 2FA codes, credential saving on login/update/signup, and secret viewing without requiring device registration).

— Names each surface and permission model but no exact menu path

Release Notes 2026 (2026-09-21)

05 **Dynamic Groups filter operators renamed** BREAKING 81

Replaces deprecated `$le` and `$ge` Dynamic Groups `memberQuery` filter operators with `$lte` and `$gte` to align with the rest of the API list endpoints; existing API scripts using `$le` or `$ge` must be updated (Admin Portal-managed Dynamic Groups are unaffected).

— Exact deprecated and replacement operator names given Release Notes 2024 (2026-09-21)

06 **Step-up authentication and MFA proof in SSO** NEW 80

Adds a `sso_auth` log event specifically for forced re-authentication events in SAML/OIDC flows for compliance visibility, `ForceAuthn=true` (SAML) and `prompt=login` (OIDC) support so applications can trigger step-up authentication without ending the user's existing session, and Authentication Methods References (AMR) in SAML assertions and OIDC ID tokens to supply MFA proof to connected applications.

Audit which authentication factors were used during an SSO login by inspecting the AMR claim in a JumpCloud-issued OIDC ID token — useful for verifying Zero Trust policy compliance.

📍 In the Admin Portal, go to Directory Insights and filter by event type `sso_auth` to view forced re-authentication events and confirm AMR tags (e.g., password + OTP) for each login.

— Names exact SAML/OIDC parameters and log event with an audit example

Release Notes 2025 (2026-09-21)

07 **LDAP security hardening** BREAKING 80

Restricts JumpCloud LDAP TLS to five modern cipher suites — `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` — removing CBC-mode and SHA-1/SHA-256 suites; any LDAP client configured to use only removed suites will fail to connect. Also adds source IP tracking for incoming LDAP requests for enhanced security and traceability.

— Names every allowed cipher suite and the connection-failure impact

Release Notes 2026 (2026-09-21) | Release Notes 2025 (2026-09-21)

---

08 **Password Manager desktop app capacity, sharing, and session controls**

IMPROVED — 78

Adds an Auto Lock setting that locks the app automatically on device lock or after a configurable auto-lock schedule (default 15 minutes of inactivity, up to 1440 minutes); raises item storage to up to 13,000 items with autofill enabled, 5,000 items per shared folder, and 30,000 items without autofill; introduces a paginated items list showing 100 items per page; adds a Move All items button for shared folders and the ability to delete multiple folders at once; adds JumpCloud Password Manager Cloud Backups as a secure fallback to minimize vault data loss; lets admins create and update shared folders and assign users or user groups to any folder, building on earlier Shared Folders management with per-folder manual sync; adds item archiving (view, delete, restore), multi-select and sorting by nickname, a Copy button in Password History, syncing of shared folder events to Directory Insights, and a What's New link plus Download Logs and Keyboard Shortcuts options under Settings.

— Names exact settings and limits but no CLI/API commands

What's New: Password Manager Desktop Ap… | Release Notes 2023 (2026-09-21)

---

09 **JumpCloud Go passwordless authentication** NEW — 78

Adds the `JumpCloud Go` browser extension for hardware-protected, phishing-resistant passwordless login on managed macOS and Windows using Apple Touch ID and Windows Hello, set as the default MFA method for conditional access policies on SSO apps with user and device verification on every session. Adds AWS Verified Access integration supporting JumpCloud Go, SAML, and OIDC SSO as a VPN-less remote access route to corporate web resources, a `User Portal > Security > Multi-Factor Authentication` section for managing Go sessions, and automatic session refresh when a user logs in to or unlocks their managed device.

— Names the extension, integration and portal path but no CLI command

Release Notes 2024 (2026-09-21) | Release Notes 2023 (2026-09-21)

---

10 **Windows MDM CSP policy expansion** NEW — 77

Adds six new Windows MDM CSP-based policies — Windows Wi-Fi Configuration Policy, Windows VPN Policy, Install Certificate Policy, Bluetooth Configuration Policy, Control Policy Conflict, and Windows Notification Service — plus a Create a Windows SCEP Profiles Policy for Simple Certificate Enrollment Protocol, and a generic Windows MDM CSP framework letting admins configure any company policy via a Configuration Service Provider. Later adds three Windows Personalization policies (Taskbar Customization, Screen Saver Configuration, Desktop and Lock Screen Settings) and Windows 11 25H2 support.

— Names every policy but no exact config path to enable them

Release Notes 2025 (2026-09-21) Release Notes 2024 (2026-09-21)

11 **JumpCloud Protect push and enrollment controls** IMPROVED

75

Adds Push Bombing Mitigation to JumpCloud Protect, blocking more than one MFA push per resource within a 60-second window (limit adjustable or disabled in MFA Configurations); enables MFA by default for all new admin accounts; adds lock-screen push notification approval and denial (long-press on iOS, expand on Android) when biometric User Verification isn't required; and adds User Group Exclusion from Unified MFA Push enforcement, letting excluded groups use TOTP only.

Enable the Push Bombing Mitigation limit or disable it entirely for high-volume MFA environments where legitimate users may trigger multiple pushes.

📍 In Admin Portal, go to Security › MFA Configurations, then adjust Maximum concurrent push attempts or toggle Push Bombing Mitigation off.

— Names the exact push limit window and a config example Release Notes 2025 (2026-09-21)

Release Notes 2023 (2026-09-21)

12 **Dynamic Device and User Groups** NEW

75

Adds Dynamic Device Groups and Dynamic User Groups with attribute-based rules that automate group membership, with an optional review step before changes take effect. Adds Vendor and Public IP Address as new attributes for dynamic device group rules, and Manager as a new attribute for dynamic user group rules.

Set up a dynamic device group that automatically includes all devices with a known public IP range, removing manual membership updates.

📍 In Admin Portal, go to Device Groups › New Group, enable Dynamic, and add a rule using the Public IP Address attribute with your target value.

— Names all three new attributes with a rule-building example

Release Notes 2023 (2026-09-21)

13 **Access Risk Detection (GA)** NEW 75

Promotes Access Risk Detection to general availability, with behavior-based detection (unusual locations, impossible travel, unfamiliar devices/browsers, dormant accounts, abnormal login volumes), automatic MFA-based resolution, intelligent suppression for VPN/gateway logins, customizable risk factor thresholds, and email/Slack/webhook notifications.

— Enumerates detection signals and notification channels, no config steps

Release Notes 2026 (2026-09-21)

14 **Password Manager desktop app platform and region expansion**

IMPROVED 73

Adds a Debian (Ubuntu x86) build and later an RPM build for Linux (Red Hat/Fedora x64) with automatic updates, a native ARM build for macOS M-series (M1+) devices, EU region support selectable on the activation page and browser-extension pairing page, and installation/update via a new command template v.2.0.0 for Windows and macOS that removes the need for ongoing admin permissions after the initial run.

Onboard EU-region users so their vault data stays within the EU — select the region before activating to ensure data residency compliance.

> 📍 On the Password Manager activation page, select `EU` as your region before entering your activation code, or on the browser extension pairing page, choose the EU region from the region selector.

— Names builds and template version with a region-selection example

What's New: Password Manager Desktop Ap… Release Notes 2023 (2026-09-21)

15 **IPA and iOS/iPadOS BYOD enrollment enhancements** NEW 73

Supports BYOD enrollments of iOS/iPadOS 18+ devices via Account-Driven User Enrollment (ADUE), now generally available; adds two policies to control Apple Intelligence features on iOS/iPadOS 18.1 and macOS 15.1 Sequoia; adds Apple Business Manager as a new service discovery method for Account-driven Enrollment; and adds support for uploading and deploying IPA-based app packages to iOS and iPadOS devices via the JumpCloud Private Repository.

— Names each Apple MDM capability, no step-by-step enrollment guide

Release Notes 2025 (2026-09-21) Release Notes 2024 (2026-09-21)

16 **Commands feature: install trigger and running status** NEW 73

Adds a `Command After Agent Install` command type so admins can specify enrollment commands to run immediately after the JumpCloud agent installs on new devices, and a `Running` column in the commands queue to distinguish commands waiting versus actively running.

Run a bootstrap script on every new device the moment the JumpCloud agent is installed, without waiting for the device to fully enroll.

> 📍 In the Admin Portal, go to Commands › New Command › select command type `Command After Agent Install`, then write your onboarding script and save.

— Named command type with a full setup walkthrough `Release Notes 2022 (2026-09-21)`

---

17 **Remote Assist remote desktop support** NEW 72

Introduces JumpCloud Remote Assist for remote desktop access and support on Windows and macOS, available to all customers on JumpCloud Agent v1.61.0+ with no limits on session time or count. Adds support for Windows Server initial login screens (entering credentials and interacting with the login window), a `Ctrl+Alt+Del` shortcut from the Session Options toolbar, and user-switching between accounts during sessions on Windows. Later ships as two architecture-specific packages (ARM64 for Apple Silicon, x86_64 for Intel Macs) and automatically detects and terminates existing unresponsive agent instances before establishing a new connection.

— Full before/after across releases but no admin UI path given

`Release Notes 2024 (2026-09-21)` `Release Notes 2023 (2026-09-21)` `Release Notes 2022 (2026-09-21)`

---

18 **Private Repository app version updates** IMPROVED 72

Enables updating an existing application's binary in JumpCloud's Private Repository via `Update Version` on the application's Details tab, deploying to already-associated devices without re-creating the application or reassigning device groups.

— Names the exact UI control and its no-recreate behavior `Release Notes 2026 (2026-09-21)`

---

19 **Knowledge Base Uploads for AI Assistant** NEW 71

Adds Knowledge Base Uploads for AI Assistant, supporting PDF, TXT, DOCX, and MD files (10 MB max per file, 25 files max per org) in a private org-scoped knowledge base, with source attribution in chat responses.

— Exact file types and limits given, no upload steps `Release Notes 2026 (2026-09-21)`

---

20 **Device management policies: macOS and Linux** NEW 70

Adds a Managed Login Items policy for macOS 13 Ventura to prevent end users from disabling persistent software like the JumpCloud Agent, new software management policies to delay or block macOS Ventura upgrades on managed devices, and a Linux

Forbidden Services policy to remove support for service types not required for everyday fleet operation.

Block unintended macOS Ventura upgrades across a managed fleet until your org is ready to migrate.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy › macOS › select `Delay or Block macOS Ventura`, configure the deferral period, and bind it to the target device group.

— Names each policy with a worked deferral configuration example

Release Notes 2022 (2026-09-21)

---

21 **Multi-Tenant Portal and MSP management** NEW 70

Adds MSP Create Custom Command Templates so admins can build reusable command templates applied across tenant organizations; ability to run ad hoc commands on devices directly from the Multi-Tenant Portal and Admin Portal; user and device management in the Multi-Tenant Portal including creating users/devices, exporting lists, and bulk actions on users; an MTP Homepage widget showing aggregate reports across all tenant orgs from the last 24 hours with bulk JSON or CSV export; and a Global MFA Requirement setting under `Account > Settings` in the Multi-Tenant Portal for bulk MFA enforcement across all MSP admins.

— Names each MTP surface and export formats, no exact click path

Release Notes 2025 (2026-09-21) Release Notes 2023 (2026-09-21)

Release Notes 2022 (2026-09-21)

---

22 **Asset Management field and cloning enhancements** IMPROVED 68

Adds a File Upload custom field type and an Invoice file-upload field to Asset Management asset records, a Disk Size field under Hardware Specifications and Local Device Accounts under Device Management on the Asset Detail page, and asset cloning to streamline creating assets with shared hardware configurations, automatically clearing unique fields to prevent duplicates.

— Names each new field and the cloning behavior Release Notes 2025 (2026-09-21)

---

23 **SCIM group management mapping UI** NEW 68

Expands SCIM group management with a dedicated mapping UI supporting name overrides, extended schema mapping, side-by-side object transformation previews, flexible custom validation with selectable test users and groups, and automatic sync logging to Directory Insights.

— Enumerates mapping capabilities, no field-level example Release Notes 2026 (2026-09-21)

24 **Device Monitoring and Alerting** NEW 67

Introduces JumpCloud Device Monitoring and Alerting, enabling customizable alerts for device health (availability, software, uptime, disk, battery), policy/command/software management results, user addition and privilege elevation directory events, and command execution monitoring across Windows, macOS, and Linux endpoints.

— Enumerates monitored signals but no alert-setup steps Release Notes 2025 (2026-09-21)

25 **Admin role management: Asset Manager and Custom Admin Roles** NEW 67

Adds an Asset Manager built-in role in the Admin Portal giving dedicated access to Asset Management without broader Admin Portal access, and Custom Admin Roles for direct organizations enabling granular role configuration based on specific responsibilities.

Grant a procurement team member access to Asset Management without exposing other Admin Portal capabilities.

> 📍 In the Admin Portal, go to Administrators › Roles, assign the built-in Asset Manager role to the user, then confirm the user can access Asset Management but not other admin areas.

— Names the built-in role with a walkthrough assigning it Release Notes 2025 (2026-09-21)

26 **RADIUS and EAP authentication enhancements** IMPROVED 67

Supports certificate (passwordless) based authentication for JumpCloud RADIUS accepting certificates from any X.509-compliant certificate authority; adds RADIUS authentication using Azure AD as the identity provider; supports multiple IP addresses on a single RADIUS server configuration; and adds faster reconnect for EAP connections by letting recently authenticated devices skip the full security handshake using a secure cached session credential.

— Names auth mechanisms across releases, no CLI setup given

Release Notes 2026 (2026-09-21) Release Notes 2025 (2026-09-21) Release Notes 2022 (2026-09-21)

27 **Shadow AI Dashboard** NEW 66

Adds a Shadow AI Dashboard — a dedicated UI for managing generative AI app usage — with insights into AI adoption, usage by department, and most active users, under the renamed AI & SaaS Management product.

Detect unauthorized or unmanaged generative AI app usage across your org before it becomes a shadow IT risk.

> 📍 In the Admin Portal, go to AI & SaaS Management › Shadow AI Dashboard to review AI adoption by department, set approval or restriction policies on discovered AI apps, and identify most active users.

— Named UI with a navigation example, limited mechanism detail

Release Notes 2025 (2026-09-21)

---

28 **API filter escaping for pipe characters** IMPROVED 65

Adds support for escaping the pipe `|` character with a backslash `\` inside `in` and `nin` API filter search terms, e.g. `GET /systemusers?filter=department:$in:...`.

— Names exact endpoint and escape syntax, no broader use case

Release Notes 2024 (2026-09-21)

---

29 **Android device management: zero-touch, EMM, and Android 15** NEW 65

Adds Android zero-touch enrollment for fully managed, dedicated, and company-owned devices with a work profile; Android EMM support for fully managed and dedicated devices with 7 new policies including Factory Reset Protection, Kiosk Mode, and System Updates; and full Android 15 MDM support.

— Names enrollment modes and example policies, no config path

Release Notes 2024 (2026-09-21) Release Notes 2023 (2026-09-21)

---

30 **SaaS Management app discovery and governance** NEW 64

Introduces SaaS Management: discover and track SaaS app usage via the JumpCloud Go browser extension, categorize apps, restrict access with warning or blocking messaging, and manually track licensing. Later enables app discovery without requiring JumpCloud Go to be enabled as a feature (the browser extension is still required), and lets admins manually add applications for tracking, warning, and blocking before automatic discovery occurs.

— Explains discovery mechanism but no navigation path Release Notes 2024 (2026-09-21)

---

31 **RSS feeds for release and bug fix announcements** NEW 64

Adds RSS feed subscriptions for JumpCloud Help Center content, with dedicated feeds for product release notes, agent release notes, and bug fix reports via Subscribe to Support RSS Feed, Subscribe to Release Notes and Bug Fixes RSS Feed, Subscribe to Agent Release Notes RSS Feed, and Subscribe to Bug Fixes RSS Feed.

Subscribe to only the product and agent release notes RSS feed to track new JumpCloud capabilities in your RSS reader.

📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click Subscribe to Release Notes and Bug Fixes RSS Feed. 3. Open your RSS reader and add the returned feed URL to begin receiving product release notes and agent release notes automatically.

— Names each feed with a subscription walkthrough

Subscribe to New Feature and Product Up…

32 **Windows App Catalog** NEW 62

Introduces the JumpCloud App Catalog for Windows with 10 apps — 1Password, Adobe Acrobat Reader DC, Chrome, Firefox, Google Drive, OneDrive, OpenVPN, Slack, Webex, and Zoom — with auto-update support.

— Lists all ten apps but no enablement steps Release Notes 2024 (2026-09-21)

33 **Browser and OS patch management reporting** NEW 60

Adds a universal Google Chrome Browser Patch Management policy covering macOS, Windows, and Linux from a single policy, an OS Patch Management Policy Report showing which policies are applied to each device and their compliance status, and a Browser Patch Management Policy Report showing devices' current browser versions for compliance audits.

— Names the cross-platform policy and both reports, no export detail

Release Notes 2022 (2026-09-21)

THINNER COVERAGE BELOW

34 **Access Requests** NEW 59

Adds JumpCloud Access Requests with configurable approval flows, request management in the Admin Portal, tracking via Directory Insights, and webhook integration with ticketing systems such as JIRA.

— Names the ticketing integration but no request-creation steps

Release Notes 2025 (2026-09-21)

35 **Conditional Access disk encryption condition** NEW 59

Adds a new device condition in Conditional Access Policies allowing or denying access based on whether disk encryption is enabled or disabled (FileVault on macOS, system drive on Windows, root disk on Linux).

— Names the platform-specific encryption checks, no policy-build steps

Release Notes 2022 (2026-09-21)

36 **MDM certificate and token expiration visibility** NEW 58

Adds the MDM Upcoming Certificate Expirations homepage widget for at-a-glance lifecycle visibility into MDM, Apple VPP, and Automated Device Enrollment certificates, and an MDM Certificate & Token Status Board in the Admin Portal and Multi-Tenant Portal providing expiry status for MDM, SSO, and AzureAD tokens and certificates with an aggregate MSP view.

— Names both surfaces tracked, no remediation steps Release Notes 2023 (2026-09-21) Release Notes 2022 (2026-09-21)

37 **Admin Portal quick search shortcut** NEW 55

Adds a `Cmd/Ctrl + K` keyboard shortcut in the Admin Portal for fast search and navigation between areas.

— Exact shortcut given, minimal further detail needed Release Notes 2024 (2026-09-21)

38 **SentinelOne agent deployment via Commands** NEW 54

Adds SentinelOne integration via the Commands feature, enabling download and installation of the SentinelOne Agent on macOS, Windows, and Linux devices.

— Names the vendor and platforms, no command syntax given Release Notes 2022 (2026-09-21)

39 **Directory Insights filter granularity** IMPROVED 54

Expands Directory Insights Application Name and Resource Name filters into sub-categories: Application Display Name, Application Display Label, Resource Display Name, and Resource Display Label.

— Names all four new sub-filters, no query example Release Notes 2025 (2026-09-21)

40 **AI Gateway third-party integrations** NEW 53

Adds AI Gateway integration support for Miro, Mixpanel, Netlify, Supabase, Vercel, and Zoho, plus new user-facing AI Gateway connectors for Antigravity and Claude Code.

— Names all eight integrations, no setup mechanism given product docs

41 **Schedule macOS software updates per device** NEW 53

Adds ability to schedule macOS software updates for a single device from the Admin Portal, including configurable install timeframe and user deferral options.

— Describes scheduling options without exact navigation path

Release Notes 2022 (2026-09-21)

42 **Background Tools remote troubleshooting** NEW 52

Adds Background Tools remote command line and remote file transfer capabilities, letting admins access and troubleshoot end-user devices without interrupting active sessions.

— Names the capability without exact invocation steps Release Notes 2023 (2026-09-21)

43 **New MFA methods: TOTP for LDAP and Duo Universal Prompt** NEW

50

Adds Time-based One-time Password (TOTP) as a supported MFA factor for LDAP applications alongside JumpCloud Protect push notifications, and adds Duo Universal Prompt support for JumpCloud MFA ahead of Duo's March 30, 2024 deprecation of the traditional prompt.

— Names both methods and Duo's deprecation date, no setup steps

Release Notes 2023 (2026-09-21) Release Notes 2022 (2026-09-21)

---

44 **JumpCloud MCP Server (GA)** NEW

50

Adds JumpCloud MCP Server (GA) to connect compatible AI tools to JumpCloud using the Model Context Protocol, enabling interaction with org data and admin capabilities from AI tools and code editors.

— Describes purpose but no connection steps or tool list Release Notes 2025 (2026-09-21)

---

45 **Self-Service Account Provisioning** NEW

50

Adds Self-Service Account Provisioning allowing new users to bind their account to JumpCloud-managed macOS and Windows devices from the login window, including wireless connectivity controls.

— Describes the login-window flow without setup steps Release Notes 2023 (2026-09-21)

---

46 **Support case portal** IMPROVED

50

Adds the JumpCloud Case Portal to centralize support cases and ideas from all administrators in one filterable, searchable view; later renamed/expanded into a Customer Portal with sortable columns for Last Modified Date, Account, and Requester, plus attachment uploads and inline ticket editing.

— Names sortable columns, no ticket-creation steps Release Notes 2025 (2026-09-21)

Release Notes 2022 (2026-09-21)

---

47 **Password Manager launch: encrypted sync, autofill, and sharing**

NEW

48

Launches JumpCloud Password Manager with end-to-end encrypted local sync, a browser extension for autofill, secure password sharing, and password generation — no master password required.

— Describes launch features but no config keys or exact steps

Release Notes 2022 (2026-09-21)

---

48 **Password policy and self-service reset behavior** IMPROVED

47

Changes the default password policy compliance behavior so users are forced to comply on their next password change rather than immediately, and enables admins to allow users to reset passwords from the User Portal when MFA is enforced but not yet enrolled.

— Describes behavior change without the config toggle name

Release Notes 2025 (2026-09-21)

---

49 **AI Commands Builder (GA)** NEW

45

Adds AI Commands Builder (GA) to create and refine commands for Windows, macOS, and Linux devices using natural-language prompts within the Admin Portal.

— States capability without a worked prompt example Release Notes 2025 (2026-09-21)

---

50 **Insights Rules frequency-based alert triggers** NEW

45

Adds alert trigger conditions based on event frequency to Insights Rules, allowing alerts to fire only when defined thresholds are met rather than on every individual event.

— Describes threshold behavior without exact configuration UI

Release Notes 2026 (2026-09-21)

---

51 **User Portal Applications tabs** IMPROVED

43

Adds All Applications and Favorites tabs to the Applications area of the User Portal, with starring and responsive search on both tabs.

— Names the tabs and search behavior, no navigation path Release Notes 2023 (2026-09-21)

---

52 **Password Manager .appimage build discontinued** DEPRECATED

40

The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian builds.

— Clear cutover date but no migration command What's New: Password Manager Desktop Ap…

---

53 **Password Date column on Users page** IMPROVED

40

Adds Password Date (last password set date) as a viewable column on the Users page in Admin Portal.

— Names the exact column with clear location, thin scope Release Notes 2023 (2026-09-21)

---

54 **SaaS License Management** NEW

38

Adds SaaS License Management (GA) for automated license tracking, usage optimization, and cost management via SaaS application connectors.

— Names the capability without connector list or setup steps

Release Notes 2025 (2026-09-21)

55 **Service Account for APIs** NEW 35

Adds Service Account for APIs, enabling API service accounts that are not tied to individual Admin accounts.

— No creation steps or scoping detail given Release Notes 2025 (2026-09-21)

56 **ADI support for Windows Server 2022** NEW 35

Adds ADI support for Windows Server 2022 domain controllers and member servers without workarounds or reduced functionality.

— States support without configuration guidance Release Notes 2023 (2026-09-21)

57 **Identity First Login flow** NEW 33

Introduces an Identity First Login flow for the User Portal and SSO applications, splitting login into a username screen followed by a separate password screen.

— Describes UX change without rollout or config detail Release Notes 2022 (2026-09-21)

58 **OIDC single sign-on support** NEW 32

Supports OpenID Connect (OIDC) Single Sign-On so users can authenticate to OIDC-capable web applications with a single set of credentials.

— States capability without configuration detail Release Notes 2022 (2026-09-21)

59 **WebAuthn enrollment split into credential types** IMPROVED 32

Splits WebAuthn enrollment into separate credential types — Security Keys and Device authenticators — for clearer enrollment.

— No mention of where enrollment lives or migration effect

Release Notes 2023 (2026-09-21)

60 **M365/Azure AD security groups in directory integration**

IMPROVED 29

Adds M365/Azure AD security groups support to the JumpCloud Microsoft 365 directory integration.

— Bare capability statement, no sync scope detail Release Notes 2022 (2026-09-21)

61 **Legacy Password Manager CSV export for Vault migration** 25

IMPROVED

Enhances CSV export for the legacy Password Manager with new data fields to support migration into JumpCloud Vault.

— No field names or migration steps given What's New: Password Manager Desktop Ap…

62 **Deel SSO Connector** NEW 23

Adds Deel SSO Connector integration with JumpCloud.

— Bare integration name with no configuration detail Release Notes 2023 (2026-09-21)

63 **Trials recap summary** NEW 21

Adds JumpCloud Trials recap summarizing feature usage at the end of a trial.

— Bare description with no detail on delivery or contents Release Notes 2022 (2026-09-21)

64 **Linux Fedora 37 agent support** NEW 21

Adds JumpCloud agent support for Linux Fedora 37.

— Bare platform-support statement Release Notes 2022 (2026-09-21)

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; all callers must switch to `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups `memberQuery` filter operators `$le` and `$ge` are deprecated in favor of `$lte` and `$gte`; existing API scripts using `$le` or `$ge` must be updated (Admin Portal-managed Dynamic Groups are not affected).

! JumpCloud LDAP no longer accepts connections using CBC-mode or SHA-1/SHA-256 integrity cipher suites (including RSA and ECDHE key exchange variants); only the five explicitly listed GCM cipher suites are permitted. Any LDAP client configured to use only the removed suites will fail to connect.

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian builds.

Release Notes 2022 (2026-09-21) BLOG ↗ STABLE

JumpCloud 2022 adds certificate-based RADIUS auth, TOTP for LDAP MFA, Remote Assist, Password Manager, and a new Command After Agent Install type.

└──▷ HOW TO FIND IT

Run a bootstrap script on every new device the moment the JumpCloud agent is installed, without waiting for the device to fully enroll.

> 📍 In the Admin Portal, go to Commands › New Command › select command type `Command After Agent Install`, then write your onboarding script and save.

Block unintended macOS Ventura upgrades across a managed fleet until your org is ready to migrate.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy › macOS › select `Delay or Block macOS Ventura`, configure the deferral period, and bind it to the target device group.

Adds `Command After Agent Install` command type so admins can specify enrollment commands to run immediately after the JumpCloud agent is installed on new devices.

Adds a `Running` column to the commands queue to distinguish commands waiting in the queue from those actively running.

Supports certificate (passwordless) based authentication for JumpCloud RADIUS, accepting certificates from any X.509-compliant certificate authority.

Adds Time-based One-time Password (TOTP) as a supported MFA factor for LDAP applications, in addition to JumpCloud Protect push notifications.

Introduces JumpCloud Remote Assist for remote desktop access and support on Windows and macOS devices, available to all customers running JumpCloud Agent v1.61.0 or later, with no limits on time or number of sessions.

`+20 more`

`Release Notes 2023 (2026-09-21)` `BLOG ↗` `STABLE`

JumpCloud 2023 adds passwordless Go auth, Android zero-touch enrollment, remote Background Tools, dynamic groups, and Duo Universal Prompt support.

└──▷ HOW TO FIND IT

Enable the Push Bombing Mitigation limit or disable it entirely for high-volume MFA environments where legitimate users may trigger multiple pushes.

> 📍 In Admin Portal, go to Security › MFA Configurations, then adjust `Maximum concurrent push attempts` or toggle Push Bombing Mitigation off.

Set up a dynamic device group that automatically includes all devices with a known public IP range, removing manual membership updates.

> 📍 In Admin Portal, go to Device Groups › New Group, enable Dynamic, and add a rule

using the `Public IP Address` attribute with your target value.

Adds `JumpCloud Go` browser extension enabling hardware-protected, phishing-resistant passwordless login on managed macOS and Windows devices using Apple Touch ID and Windows Hello.

Adds `JumpCloud Go` as the default MFA method for conditional access policies on SSO apps, with user and device verification on every application session.

Adds `AWS Verified Access` integration supporting JumpCloud Go, SAML, and OIDC SSO as a VPN-less remote access solution for corporate web resources.

Adds Android zero-touch enrollment support for fully managed, dedicated, and company-owned Android devices with a work profile.

Adds `Background Tools` remote command line and remote file transfer capabilities, letting admins access and troubleshoot end-user devices without interrupting active sessions.

`+23 more`

`Release Notes 2024 (2026-09-21)` `BLOG ↗` `STABLE`

JumpCloud 2024 adds Device Monitoring & Alerting, Windows App Catalog, SaaS Management, BYOD iOS 18 MDM, CSP policy framework, and more.

└──▷ **TRY IT**

Reactivate a user account after migrating scripts from the deprecated GET endpoint to the new POST endpoint.

```
$ POST https://console.jumpcloud.com/systemusers/reactivate/{id}
{
  "activation_email": "{email_address}"
}
```

Adds `GET /systemusers?filter=department:$in:...` support for escaping the pipe `|` character with backslash `\` inside `in` and `nin` API filter search terms.

Adds six new Windows MDM CSP-based policies: Windows Wi-Fi Configuration Policy, Windows VPN Policy, Install Certificate Policy, Bluetooth Configuration Policy, Control Policy Conflict, and Windows Notification Service.

Adds a new Windows MDM policy to configure Simple Certificate Enrollment Protocol (`SCEP`) for Windows devices via `Create a Windows SCEP Profiles Policy` .

Replaces the deprecated `GET api/systemusers/reactivate/{id}` endpoint with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` accepting `activation_email` in the request body.

Replaces deprecated `$le` and `$ge` Dynamic Groups `memberQuery` filter operators with `$lte` and `$gte` to align with the rest of the API list endpoints.

`+18 more`

└─▷ **BREAKING ON UPGRADE**

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; all callers must switch to `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! Dynamic Groups `memberQuery` filter operators `$le` and `$ge` are deprecated in favor of `$lte` and `$gte`; existing API scripts using `$le` or `$ge` must be updated (Admin Portal-managed Dynamic Groups are not affected).

`Release Notes 2025 (2026-09-21)` `BLOG ↗` `STABLE`

JumpCloud adds Access Requests, MCP Server, Shadow AI Dashboard, forced re-auth/AMR claims, Custom Admin Roles, and IPA file deployment in 2025 releases.

└─▷ **HOW TO FIND IT**

Audit which authentication factors were used during an SSO login by inspecting the AMR claim in a JumpCloud-issued OIDC ID token — useful for verifying Zero Trust policy compliance.

📍 In the Admin Portal, go to Directory Insights and filter by event type `sso_auth` to view forced re-authentication events and confirm AMR tags (e.g., password + OTP) for each login.

Detect unauthorized or unmanaged generative AI app usage across your org before it becomes a shadow IT risk.

📍 In the Admin Portal, go to AI & SaaS Management › Shadow AI Dashboard to review AI adoption by department, set approval or restriction policies on discovered AI apps, and identify most active users.

Grant a procurement team member access to Asset Management without exposing other Admin Portal capabilities.

📍 In the Admin Portal, go to Administrators › Roles, assign the built-in `Asset Manager` role to the user, then confirm the user can access Asset Management but not other admin areas.

Adds `sso_auth` event to logs specifically for forced re-authentication events in SAML/OIDC flows, providing visibility for compliance and security reviews.

Adds JumpCloud MCP Server (GA) to connect compatible AI tools to JumpCloud using the Model Context Protocol, enabling interaction with org data and admin capabilities from AI tools and code editors.

Adds JumpCloud Access Requests with configurable approval flows, request management in the Admin Portal, tracking via Directory Insights, and webhook integration with ticketing systems such as JIRA.

Adds Shadow AI Dashboard — a dedicated UI for managing generative AI app usage, with insights into AI adoption, usage by department, and most active users — under the renamed AI & SaaS Management product.

Adds `Asset Manager` role in the Admin Portal, providing dedicated access to Asset Management without granting broader Admin Portal access.

`+19 more`

└─▷ **BREAKING ON UPGRADE**

! JumpCloud LDAP no longer accepts connections using CBC-mode or SHA-1/SHA-256 integrity cipher suites (including RSA and ECDHE key exchange variants); only the five explicitly listed GCM cipher suites are permitted. Any LDAP client configured to use only the removed suites will fail to connect.

`Release Notes 2026 (2026-09-21)` **BLOG ↗** STABLE

JumpCloud adds Password Vault across Admin/User Portal, mobile, and browser extension, plus GA Access Risk Detection and Agent Identities.

└─▷ **HOW TO FIND IT**

Register a new AI agent and scope its OAuth credentials to require credentials in the request body — needed when integrating with platforms like Amazon Bedrock AgentCore that enforce `client_secret_post`.

📍 In the Admin Portal: Go to Identity Management › Agents › + Add. Fill in agent name, owner, and company email on the Details tab. On the Authentication tab, generate OAuth client credentials, then set the client authentication type to `Client secret post` to send the client ID and secret in the request body.

Surface which SaaS apps have locally running MCP servers and connect them to AI Gateway for centralized governance.

📍 In the Admin Portal: Go to Access › AI & SaaS Management › Applications. Review the MCP Connection column for apps showing `Discovered`. Click the status for a target app, review discovered MCP servers, then click `Add MCP Server` to register them in AI Gateway.

Adds `MCP Connection` column (showing `Discovered`, `Connected`, or `Not Available`) in `Access > AI & SaaS Management > Applications` to surface MCP server status and enable adding servers to AI Gateway without switching tools.

Launches Password Vault in the JumpCloud Protect mobile app, letting users view secrets, autofill credentials into mobile apps, and autofill in the mobile browser using Push/TOTP-authenticated sessions.

Launches Password Vault in the Admin Portal, enabling admins to enable the vault, assign access to users and user groups, manage shared credentials and folders with Connect/View/Edit/Manage permissions, and view audit logs in Directory Insights.

Launches Password Vault in the User Portal, letting users manage and share credentials and website applications, bulk import via CSV or Excel, and group items in personal and shared folders.

Launches Password Vault Browser Extension via JumpCloud Go, supporting autofill of Non-SSO website passwords and 2FA codes, credential saving on login/update/signup, and secret viewing — without requiring device registration.

`+8 more`

`Subscribe to New Feature and Product Update Announcements (2026-09-21)` **BLOG ↗**

STABLE

JumpCloud adds RSS feed subscriptions for release notes, product updates, agent notes, and bug fix announcements.

└──▷ HOW TO FIND IT

Subscribe to only the product and agent release notes RSS feed to track new JumpCloud capabilities in your RSS reader.

📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click `Subscribe to Release Notes and Bug Fixes RSS Feed`. 3. Open your RSS reader and add the returned feed URL to begin receiving product release notes and agent release notes automatically.

`What's New: Password Manager Desktop App (2026-09-21)` **BLOG ↗** **LATEST**

STABLE

JumpCloud Password Manager Desktop App adds EU region support, raised item limits, auto-lock scheduling, ARM build for Apple Silicon, and more.

└──▷ HOW TO FIND IT

Onboard EU-region users so their vault data stays within the EU — select the region before activating to ensure data residency compliance.

📍 On the Password Manager activation page, select `EU` as your region before entering your activation code, or on the browser extension pairing page, choose the EU region from the region selector.

Introduces an auto-lock schedule with a default of 15 minutes of inactivity and a configurable maximum of 1440 minutes (24 hours).

Adds `Auto Lock` setting that locks the Password Manager app automatically when the device is locked.

Scales item storage capacity to up to 13,000 items with autofill enabled, up to 5,000 items per shared folder, and up to 30,000 items when autofill is not used.

Adds a `Move All` items button in shared folders to move all items from a folder in a single action.

Introduces paginated items list displaying 100 items per page.

`+14 more`

└─▷ **BREAKING ON UPGRADE**

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian builds.

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ announced · Sep 18, 2026 BLOG ↗ RANK

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0 shipped three Early Access/beta capabilities this window: an FGA Permissions Index for constant-time permission checks, an Agent Gateway for governing AI agent tool calls, and an Identity Conversion Suite adding Anonymous Sessions and auth-flow A/B testing.

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

01 Anonymous Sessions for guest identity tracking NEW

85

Issues OAuth 2.0-compliant access tokens to unauthenticated guest users, enabling guest checkout, cart tracking, pre-login personalization, and secure API access without authentication. The anonymous session token is passed to the Auth0 `/authorize` endpoint on sign-up or login, making all prior anonymous activity available inside Actions for custom handling, with anonymous metadata transferred seamlessly into the authenticated profile. Available in Early Access to Auth0 Enterprise plan customers.

Trigger an anonymous session and obtain an OAuth 2.0 access token for a guest user before they authenticate, so cart and wishlist metadata is tracked from first interaction.

> 📍 Pass referral IDs, shopping cart IDs, or wishlist URLs to Auth0 at the start of the guest session. Auth0 issues a session token and an access token; use the access token to authorize API calls normally. When the guest signs up or logs in, include the session token in the request to the `/authorize` endpoint so Auth0 transfers all anonymous activity to the authenticated profile.

— Names `/authorize` endpoint and gives a runnable step sequence launch-20260918-8a7c8e15

02 Experiment Center for auth flow A/B testing NEW

78

Splits authentication traffic into variants (by percentage or segment) for capabilities such as passkeys, MFA steps, signup form copy, and branding themes across Universal Login (including Auth0 Custom Universal Login / ACUL), Forms, Actions, and Page Templates, without changing application code. Captures exposure, success rate, and time-to-auth per variant and segment as dedicated experiment events, enabling data-backed promotion of winning flows. Available in Early Access to Auth0 Enterprise plan customers.

A/B test a passkey-first signup flow against the standard flow for warm visitors (those with saved cart items) to measure completion rate and time-to-auth before a full rollout.

> 📍 In Experiment Center, define a feature flag for the passkey signup variant, assign the `saved cart` segment to the passkey variant and the remainder to control, then monitor the dedicated experiment events for exposure, success rate, and time-to-auth split by segment before promoting the winner tenant-wide.

— Clear UI-based starting point via feature flags and segments `launch-20260918-8a7c8e15`

---

03 Agent Gateway for AI agent tool-call authorization NEW

76

New beta control plane that evaluates tool-call authorization in real time using agent, user, and organization context before any external tool is invoked. Integrates with Auth0 Token Vault and on-behalf-of token exchange to broker short-lived, user-scoped tokens, keeping raw credentials out of LLM prompts, context windows, and logs. Supports organization-aware, multi-tenant routing that isolates tool access and connected credentials (e.g. Slack, Microsoft Teams) per customer tenant, includes an agent kill switch to disable an agent or restrict its access across a tenant, and logs every agent tool call to tenant logs with optional streaming export to SIEM services. Exposes Agent Gateway APIs for early integration during the beta, with documentation and integration guides.

— Names integrations and controls but no exact endpoint or flag `launch-20260918-5ac20484`

---

04 FGA Permissions Index for constant-time access checks NEW

68

Precomputes every access implied by chosen relations (e.g. `can_view`) on an object type and streams the results directly into the customer's own database, enabling a simple SQL JOIN instead of runtime graph traversal for permission-filtered queries. Supports indexing high-cardinality relations so lookups on indexed relations become constant-time, while non-indexed relations continue to be evaluated by FGA in real time.

— Explains mechanism and scope but gives no API or config surface

`launch-20260918-135298a6`

`announced · Sep 18, 2026` `BLOG ↗`

Auth0 FGA Permissions Index (Early Access) precomputes all user permissions and streams them into your own database for constant-time SQL JOIN-based access checks at scale.

Introduces FGA Permissions Index (Early Access), which precomputes every access implied by chosen relations (e.g. `can_view`) on an object type and streams the

results directly into the customer's own database, enabling a simple SQL JOIN instead of runtime graph traversal for permission-filtered queries.

Supports indexing high-cardinality relations on an object type so that permission lookups on indexed relations become constant-time, while non-indexed relations continue to be evaluated by FGA in real time.

announced · Sep 18, 2026 BLOG ↗

Auth0 Agent Gateway (beta) adds a CIAM control plane for governing customer-facing AI agents across multi-tenant SaaS products.

New Auth0 Agent Gateway beta provides a centralized policy-enforcement layer that evaluates tool-call authorization in real time using agent, user, and organization context before any external tool is invoked.

Integrates with Auth0 Token Vault and on-behalf-of token exchange to broker short-lived, user-scoped tokens for internal and external APIs — keeping raw credentials out of LLM prompts, context windows, and logs.

Supports organization-aware, multi-tenant routing that isolates tool access and connected credentials (e.g. Slack, Microsoft Teams) per customer tenant, preventing cross-tenant data exposure.

Provides an agent kill switch — centralized administrative controls to disable an agent or restrict its access to downstream services and tools across a tenant or deployment.

Logs every agent tool call to tenant logs with optional streaming export to SIEM services for enterprise auditability.

+1 more

announced · Sep 18, 2026 BLOG ↗

Auth0 launches Identity Conversion Suite (Early Access) with Anonymous Sessions and Experiment Center for auth flow A/B testing

└──▷ **HOW TO FIND IT**

Trigger an anonymous session and obtain an OAuth 2.0 access token for a guest user before they authenticate, so cart and wishlist metadata is tracked from first interaction.

📍 Pass referral IDs, shopping cart IDs, or wishlist URLs to Auth0 at the start of the guest session. Auth0 issues a session token and an access token; use the access token to authorize API calls normally. When the guest signs up or logs in, include the session token in the request to the `/authorize` endpoint so Auth0 transfers all anonymous activity to the authenticated profile.

A/B test a passkey-first signup flow against the standard flow for warm visitors (those with saved cart items) to measure completion rate and time-to-auth before a full rollout.

> 📍 In Experiment Center, define a feature flag for the passkey signup variant, assign the `saved cart` segment to the passkey variant and the remainder to control, then monitor the dedicated experiment events for exposure, success rate, and time-to-auth split by segment before promoting the winner tenant-wide.

Passes the anonymous session token to the Auth0 `/authorize` endpoint on sign-up or login, making all prior anonymous activity available inside Actions for custom handling.

Adds `Anonymous Sessions`: issues OAuth 2.0-compliant access tokens to unauthenticated guest users, enabling guest checkout, cart tracking, pre-login personalization, and secure API access without authentication — with seamless transfer of anonymous metadata into the authenticated profile on sign-up or login.

Adds `Experiment Center`: split authentication traffic into variants (by percentage or segment) for capabilities such as passkeys, MFA steps, signup form copy, and branding themes — across Universal Login (including Auth0 Custom Universal Login / ACUL), Forms, Actions, and Page Templates — without changing application code.

Experiment Center captures exposure, success rate, and time-to-auth per variant and segment as dedicated experiment events, enabling data-backed promotion of winning flows.

Both Anonymous Sessions and Experiment Center are available in Early Access to Auth0 Enterprise plan customers.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

BigID

Feedback ⓘ

Meet AgentIQ: Run Your Entire Data & AI Program From a Prompt (2026-09-21)

BLOG ↗ RANK

Secure your data everywhere—cloud, on-prem, and AI. BigID delivers DSPM and AI governance trusted by Fortune 500 enterprises.

> BigID launched AgentIQ, a new agentic automation layer that lets teams run data and AI security workflows via natural language, exposing discovery, classification, remediation and governance to external AI assistants through MCP, alongside pre-built agents, human-in-the-loop approval controls, shadow AI discovery, and automated reporting across 200+ integrations.

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

01 **AgentIQ agentic automation layer via MCP** NEW 70

AgentIQ lets practitioners describe data and AI security workflows in natural language and have an agent investigate, prioritize, report, and take action across BigID's discovery, classification, access intelligence, and remediation capabilities. It exposes BigID's operational capabilities — discover, classify, investigate, remediate, govern, and fulfill — to external AI assistants (Claude, ChatGPT, Copilot, Gemini, and custom enterprise agents) through `MCP` and APIs, so teams can operate BigID from whichever AI interface they already use.

— Names MCP and specific AI assistants but no exact commands or endpoints.

Meet AgentIQ: Run Your Entire Data & AI…

02 **Human-in-the-loop controls for agent actions** NEW 60

Revoking access, quarantining data, and deleting records can require human approval, with all agent actions logged and attributable to the requesting user's permissions enforced at the API and MCP layer.

— Explains approval and logging mechanism but no config path given.

Meet AgentIQ: Run Your Entire Data & AI…

THINNER COVERAGE BELOW

03 **Pre-built AgentIQ agents for common workflows** NEW 50

Ships pre-built agents for sensitive data remediation, AI governance enforcement, access exposure analysis, and privacy operations, enabling teams to start without building custom agent workflows.

— Names agent categories but no configuration detail.

Meet AgentIQ: Run Your Entire Data & AI…

04 **Shadow AI and sanctioned AI discovery workflows** NEW 45

Agents can inventory AI systems across the environment, map the data each system consumes, and flag policy violations.

— Describes function but no scope or trigger mechanism.

Meet AgentIQ: Run Your Entire Data & AI…

05 **On-demand and scheduled AgentIQ reporting** NEW 45

Supports on-demand and scheduled report generation across board-level summaries, auditor evidence packs, and recurring security team updates, driven by metadata, scan findings, and activity data.

— Lists report types but no format or scheduling detail.

Meet AgentIQ: Run Your Entire Data & AI…

06 **200+ integrations through AgentIQ framework** NEW 35

Covers more than 200 integrations spanning on-premises, cloud, SaaS, and API sources through the AgentIQ framework.

— Gives a number but no list of named integrations.

Meet AgentIQ: Run Your Entire Data & AI…

Meet AgentIQ: Run Your Entire Data & AI Program From a Prompt (2026-09-21)

BLOG ↗ LATEST STABLE

BigID launches AgentIQ, enabling agentic automation of data security workflows via MCP through Claude, Copilot, ChatGPT, Gemini, and custom enterprise agents.

Introduces AgentIQ, a new agentic automation layer that lets practitioners describe data and AI security workflows in natural language and have an agent investigate, prioritize, report, and take action across BigID's discovery, classification, access intelligence, and remediation capabilities.

Exposes BigID's operational capabilities — discover, classify, investigate, remediate, govern, and fulfill — to external AI assistants (Claude, ChatGPT, Copilot, Gemini, and custom enterprise agents) through `MCP` and APIs, so teams can operate BigID from whichever AI interface they already use.

Ships pre-built agents for sensitive data remediation, AI governance enforcement, access exposure analysis, and privacy operations, enabling teams to start without building custom agent workflows.

> Supports human-in-the-loop controls where revoking access, quarantining data, and deleting records can require human approval, with all agent actions logged and attributable to the requesting user's permissions enforced at the API and MCP layer.
>
> Enables shadow AI and sanctioned AI discovery workflows: agents can inventory AI systems across the environment, map the data each system consumes, and flag policy violations.
>
> +2 more

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.48.0 (2026-09-21) BLOG ↗ RANK

RegScale is a compliance and security governance platform that automates regulatory assessment, documentation, and continuous monitoring across organizations.

> RegScale's CLI 6.48.0 release adds finer control over JCAM import scoping and enriches Wiz issue sync with additional context fields.

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

01 **Repeatable plan ID argument for JCAM imports** NEW 60

JCAM import commands gain a repeatable `plan id` argument that lets a sync run be scoped to specific plans without editing the configured filter.

— Names exact flag but no example of use CLI 6.48.0 (2026-09-21)

THINNER COVERAGE BELOW

02 **Enriched Wiz issue sync metadata** IMPROVED 50

Wiz issue sync now imports issue notes, the affected resource's cloud console link, project business impact, policy type, and linked service tickets.

— Lists new fields but no config or command shown CLI 6.48.0 (2026-09-21)

> CLI 6.48.0 (2026-09-21) BLOG ↗ LATEST STABLE
>
> RegScale CLI 6.48.0 adds repeatable JCAM plan ID scoping and enriched Wiz issue sync with notes, links, and metadata.
>
> Adds a repeatable `plan id` argument to JCAM import commands so a sync run can be scoped to specific plans without editing the configured filter.
>
> Wiz issue sync now imports issue notes, the affected resource's cloud console link, project business impact, policy type, and linked service tickets.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ 2026_09_20_CLI_0.7.0

NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> AI SAFE² CLI 0.7.0 introduces two new `safe2 evidence` subcommands for correlating multi-harness agent evidence and detecting agent configuration drift, backed by new contract schemas and hardened evidence-handling protections.

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

01 **`safe2 evidence truth` correlation command** NEW 95

Adds `safe2 evidence truth` to correlate harness evidence, revision-bound task receipts, coverage, usage declarations, and completion claims across multiple agent harnesses. Uses the new `operational-truth-source-v1` contract schema to declare the task, revision, required evidence, receipt bindings, decision owner, assumptions, and exclusions, and the `operational-truth-manifest-v1` schema to record the correlated result, decision gate (`hold` , `review` , `ready_for_human_decision`), conflicts, gaps, and evidence boundaries. Supports `--baseline` , `--output` , `--strict` , and `--card` flags, with `--card` generating a human decision card alongside canonical JSON output from the same evaluated evidence record.

— Named command, schemas and flags with clear mechanism and behaviour.

2026_09_20_CLI_0.7.0

02 **`safe2 evidence changes` drift detection command** NEW 95

Adds `safe2 evidence changes` to run a bounded one-shot inventory of skills and recognized agent configuration locally or in CI, detecting additions, modifications, and removals. It recognizes skill packages rooted by `SKILL.md` , MCP configuration, OpenClaw configuration, and supported configuration beneath `.codex` , `.claude` , `.hermes` , and `.openclaw` , records results via the new `change-monitor-v1` contract schema without exporting configuration contents, and enforces bounded traversal of up to 10,000 entries and 1 MB per tracked file by default. Also supports the `--baseline` , `--output` , `--strict` , and `--card` flags.

— Named command, schema, recognized paths and limits, fully actionable.

2026_09_20_CLI_0.7.0

THINNER COVERAGE BELOW

03 **Strict evidence-contract protections** NEW 55

Adds strict evidence-contract protections applying to the evidence commands: duplicate-key rejection, non-finite-number rejection, symbolic-link and reparse-point guards, no-overwrite behavior, and exclusive complete-before-visible evidence publication.

— Lists distinct protections but no flags or commands to invoke them.

2026_09_20_CLI_0.7.0

2026_09_20_CLI_0.7.0 NOTES ↗ LATEST STABLE

AI SAFE² CLI 0.7.0 adds `safe2 evidence truth` and `safe2 evidence changes` to correlate multi-harness agent evidence and detect configuration drift.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2026_09_20_CLI_0.7.0
https://github.com/CyberStrategyInstitute/ai-safe2-framework.git
# already have the repo? check out this version:
$ git checkout 2026_09_20_CLI_0.7.0
```

Adds `safe2 evidence truth` command to correlate harness evidence, revision-bound task receipts, coverage, usage declarations, and completion claims across multiple agent harnesses.

Adds `safe2 evidence changes` command to run a bounded one-shot inventory of skills and recognized agent configuration locally or in CI, detecting additions, modifications, and removals.

Adds `--baseline`, `--output`, `--strict`, and `--card` flags to `safe2 evidence` subcommands for controlling baseline comparison, artifact output paths, strict-mode enforcement, and human-readable decision card generation.

Introduces `operational-truth-source-v1` contract schema to declare the task, revision, required evidence, receipt bindings, decision owner, assumptions, and exclusions.

Introduces `operational-truth-manifest-v1` contract schema to record the correlated result, decision gate (`hold`, `review`, `ready_for_human_decision`), conflicts, gaps, and evidence boundaries.

+5 more

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ v0.45.0 NOTES ↗ RANK

Docker Sandboxes isolates applications in lightweight containers, providing secure, reproducible environments for development and deployment.

> Docker Sandboxes v0.45.0 introduces v3 OCI kits with composable workload and mixin roles and kit sets, adds `sbx run --model` support for OpenAI/Anthropic-compatible endpoints, and hardens sandbox network policy enforcement, alongside a wave of new CLI flags and several breaking changes to error handling and destructive-command confirmation.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 New CLI flags across MCP, list, network and policy commands

NEW 85

Adds `sbx mcp add --disable-http2` to disable HTTP/2 for a remote MCP transport, working around servers whose HTTP/2 handling stalls long-lived streams. Adds `--quiet` (`-q`) flag to MCP server, secret, skill, template, volume, and policy-profile list commands for name-only output. Adds `--protocol` flag to local network allow rules so UDP can be selected explicitly alongside the default TCP coverage. Adds `--created-via` filter to `sbx policy ls`, which also now shows how each rule was created.

Allow outbound UDP (e.g. for a DNS-over-UDP workflow) in a sandbox network policy instead of the TCP-only default.

```
$ sbx policy allow network --protocol udp 8.8.8.8:53
```

Work around a remote MCP server that stalls long-lived HTTP/2 streams by forcing HTTP/1.1.

```
$ sbx mcp add --disable-http2 my-mcp-server
https://mcp.example.com/sse
```

— Every flag named verbatim with two runnable examples v0.45.0

02 Breaking CLI behavior changes for removals and errors

83

BREAKING

`sbx mcp catalog` has been removed; register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead. `sbx secret rm` now

returns an error on stderr when the requested secret does not exist (previously silent). `sbx mcp rm` now returns an error when the requested MCP server is not registered (previously silent). Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined; use `--force` or `--yes` / `-y` in non-interactive workflows. `sbx exec -d` / `--detach` now fails immediately for unsupported detached execution instead of silently running in the foreground.

— Five named commands with exact before/after behaviour and workaround flags v0.45.0

---

03 **Model provider support for `sbx run --model`** NEW 80

`sbx run --model` now supports any OpenAI- or Anthropic-compatible endpoint configured via the new `model.providers` setting.

Route agent runs through a custom Anthropic-compatible endpoint instead of the built-in provider.

```
$ sbx run --model my-anthropic-proxy/claude-3-5-sonnet 'summarize the repo'
```

— Named setting plus a runnable example command v0.45.0

---

04 **`sbx kit validate` rejects malformed API-key declarations** NEW 78

`sbx kit validate` now rejects malformed API-key declarations, including invalid names, missing injection domains, invalid format placeholders, and Basic-auth usernames containing a colon, and warns on no-op or out-of-allowlist declarations.

— Specific validation cases named, no example run v0.45.0

---

05 **v3 OCI kits with workload/mixin composition and kit sets** NEW 73

v3 kits are OCI-based packages with separate workload and mixin roles that combine base environments, tools, configuration, credentials, network access, and agent instructions. Kit sets let authors combine a workload and mixins, pin component versions, and publish the result as a single OCI reference with its own capabilities, lifecycle hooks, instructions, and arguments. V3 kits can scope network access by HTTP method and path, declare install-phase network access and credential use, and specify where an agent reads shared skills. Multiple OAuth-backed agents can now be composed in the same sandbox with credentials scoped to the individual kits that request them. V3 workloads and mixins must be used together; they cannot be combined with v1 or v2 kits.

— Rich mechanism but no exact command shown to try it v0.45.0

---

06 **OpenCode kit model variants and Copilot integration** NEW

70

`sbx run opencode --model` now exposes the model's supported thinking levels as OpenCode variants, selectable with Ctrl+T. The OpenCode kit now configures GitHub Copilot from the account's stored GitHub credential, enabling Copilot models without a separate device login.

— Command and keybinding named, mechanism explained v0.45.0

---

07 **Expanded JSON output for sandbox listing and inspection**

IMPROVED

60

Adds `created_at`, CPU limits, and memory limits fields to `sbx ls --json` and `sbx inspect --json` output.

— Exact fields and commands named but no example shown v0.45.0

THINNER COVERAGE BELOW

08 **Malformed network pattern validation in policy commands**

IMPROVED

55

`sbx policy allow network`, `sbx policy deny network`, `--allow-network`, and `--deny-network` now reject malformed patterns before saving them.

— Commands named but behaviour described in one line v0.45.0

---

09 **Network policy enforcement hardening** IMPROVED

55

Reverse-DNS lookups are now restricted to destination IPs already authorized by policy, closing a previous policy bypass. DNS resolution is blocked when no network rule permits it. Experimental outbound UDP now follows sandbox network policy.

— Clear before/after but no commands to act on v0.45.0

---

10 **macOS distribution as signed Sbx.app bundle** IMPROVED

50

macOS distributions now ship as a single signed `Sbx.app` bundle, with Homebrew and tarball PATH installs working via a symlink into the bundle.

— Packaging change described but no install steps given v0.45.0

---

11 **Live update of global service secrets to running sandboxes**

IMPROVED

40

Adding, updating, or removing global service secrets now updates existing local sandboxes without a restart.

— Behaviour stated plainly with no named command v0.45.0

---

12 **OAuth credential refresh coordination across sandboxes**

IMPROVED

35

OAuth refreshes are coordinated across sandboxes that share credentials, preventing simultaneous refreshes from forcing another sign-in.

— Describes benefit but no mechanism detail or surface v0.45.0

└──▷ BREAKING ON UPGRADE

! `sbx mcp catalog` has been removed; register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead.

! `sbx secret rm` now returns an error on stderr when the requested secret does not exist (previously silent).

! `sbx mcp rm` now returns an error when the requested MCP server is not registered (previously silent).

! V3 workloads and mixins must be used together; they cannot be combined with v1 or v2 kits.

! Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined; use `--force` or `--yes` / `-y` in non-interactive workflows.

! `sbx exec -d` / `--detach` now fails immediately for unsupported detached execution instead of silently running in the foreground.

v0.45.0 NOTES ↗ LATEST STABLE

Docker Sandboxes v0.45.0 ships v3 OCI kits with workload/mixin composition, new model provider settings, and a raft of new CLI flags.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.45.0 https://github.com/docker/sbx-
releases.git
# already have the repo? check out this version:
$ git checkout v0.45.0
```

└──▷ TRY IT

Route agent runs through a custom Anthropic-compatible endpoint instead of the built-in provider.

```
$ sbx run --model my-anthropic-proxy/claude-3-5-sonnet 'summarize
the repo'
```

Allow outbound UDP (e.g. for a DNS-over-UDP workflow) in a sandbox network policy instead of the TCP-only default.

```
$ sbx policy allow network --protocol udp 8.8.8.8:53
```

Work around a remote MCP server that stalls long-lived HTTP/2 streams by forcing HTTP/1.1.

```
$ sbx mcp add --disable-http2 my-mcp-server
https://mcp.example.com/sse
```

Adds `sbx run --model` support for any OpenAI- or Anthropic-compatible endpoint configured via the new `model.providers` setting.

Adds `sbx mcp add --disable-http2` to disable HTTP/2 for a remote MCP transport, working around servers whose HTTP/2 handling stalls long-lived streams.

Adds `--quiet` (`-q`) flag to MCP server, secret, skill, template, volume, and policy-profile list commands for name-only output.

Adds `--protocol` flag to local network allow rules so UDP can be selected explicitly alongside the default TCP coverage.

Adds `--created-via` filter to `sbx policy ls`, which also now shows how each rule was created.

+16 more

└─▷ **BREAKING ON UPGRADE**

- ! `sbx mcp catalog` has been removed; register a remote MCP server with `sbx mcp add` and then authorize it with `sbx mcp auth` instead.
- ! `sbx secret rm` now returns an error on stderr when the requested secret does not exist (previously silent).
- ! `sbx mcp rm` now returns an error when the requested MCP server is not registered (previously silent).
- ! V3 workloads and mixins must be used together; they cannot be combined with v1 or v2 kits.
- ! Commands that remove resources now prompt for confirmation and return a non-zero exit code when the prompt is declined; use `--force` or `--yes` / `-y` in non-interactive workflows.
- ! `sbx exec -d` / `--detach` now fails immediately for unsupported detached execution instead of silently running in the foreground.

WAS THIS USEFUL? 👍 👎

◆ Internet Services

Cloudflare

Feedback ⓘ

WAF - Enforce positive security with Application Profiles (2026-09-07) BLOG ↗

RANK

Connect, protect, and build everywhere.

> Cloudflare's biggest shift this window is Application Profiles, a new positive-security enforcement layer for WAF, alongside a sweeping expansion of the Workers AI model catalog and general availability for WAN Unified Routing; smaller improvements landed for Workflows, Hyperdrive, RPC tracing, and Browser Rendering debugging.

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

01 **New models added to Workers AI catalog** NEW 87

Adds 15 new models to Workers AI: `wan-3.0-prime` (Alibaba text-to-video, 480P/720P/1080P, adaptive aspect ratio, configurable duration); `h3` (multimodal video generation, text-to-video, first/last frame image-to-video, reference-to-video, 768P/2K); `h3-max` (fast multimodal video generation, text-to-video, image-to-video, 480P/768P); `gemini-omni-1.1-flash` and `gemini-omni-flash` (multimodal video generation/editing with conversational controls and generated audio); `gemini-3.8-flash` (long-horizon software engineering, autonomous agents, enterprise workflows); `gpt-6-astra` (OpenAI, complex reasoning, coding, computer use, research, document creation); `gpt-image-2.5-flare` and `gpt-image-2.5-sunburst` (image generation/editing with `low`, `medium`, `high`, `xhigh`, `max`, `auto` quality settings); `gpt-oss-120b` and `gpt-oss-20b` (open-weight models for production and low-latency/local use respectively); `union-alpha` (multimodal, research/coding/agentic workflows); `jev` (TypeSafe structured evaluation, typed Noul, Choice, and Score questions with calibrated probabilities and confidence scores); `pareto` (Unbiased, runs multiple language models in parallel per request and synthesizes one response, text and vision inputs); and `FLUX.2 [klein] 9B` (9-billion-parameter image generation with multi-reference editing).

Generate a 1080P cinematic video from a text prompt using the new wan-3.0-prime model via the Workers AI REST API.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/<ACCOUNT_ID>/ai/run
/@cf/alibaba/wan-3.0-prime -H 'Authorization: Bearer <API_TOKEN>'
-H 'Content-Type: application/json' -d '{"prompt": "A drone shot
of a mountain range at sunrise", "resolution": "1080P"}'
```

Evaluate a model output against structured Choice and Score questions using the jev structured evaluation model.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/<ACCOUNT_ID>/ai/run
/@cf/typesafe/jev -H 'Authorization: Bearer <API_TOKEN>' -H
'Content-Type: application/json' -d '{"state": "The response was
accurate and concise.", "questions": [{"type": "Choice", "text":
"Is the response helpful?"}, {"type": "Score", "text": "Rate
response quality 1-5"}]}'
```

Generate a high-quality image with gpt-image-2.5-sunburst using the `xhigh` quality setting for detailed output.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/<ACCOUNT_ID>/ai/run
/@cf/openai/gpt-image-2.5-sunburst -H 'Authorization: Bearer
<API_TOKEN>' -H 'Content-Type: application/json' -d '{"prompt":
"A photorealistic castle on a cliff at dusk", "quality":
"xhigh"}'
```

— Names every model and spec, with runnable curl examples for some product docs

02 **Application Profiles positive-security enforcement in WAF**

NEW 83

Introduces Application Profiles, a new layer for Cloudflare WAF that learns expected request structure and classifies traffic as conforming or non-conforming without blocking. The first profile type, Schema Profiles, models path variables, query parameters, headers, cookies, JSON bodies, and form-encoded bodies — including field types, numeric ranges, string lengths, and character classes. A new Profile Analysis view in Security Analytics lets teams review conformance trends and sampled violation details before enforcing, and profiles can then be enforced via Custom Rules scoped by hostname, path, operation, or other security signals such as Attack Score.

Review Schema Profile conformance trends and violation samples in Security Analytics before committing to enforcement.

📍 In the Cloudflare dashboard, go to Security › Analytics › Profile Analysis to inspect conformance trends and sampled non-conforming requests for your Schema Profile.

— Names mechanism, sub-features, and exact dashboard path

WAF - Enforce positive security with Ap…

03 **Batch deletion of Workflow instances** NEW 65

Enables deleting one or up to 100 Workflow instances and their stored state in a single operation via a binding or Wrangler.

— Names the limit and both entry points, no exact command product docs

04 **Hyperdrive support for Python Workers** NEW 65

Adds Hyperdrive support for Python Workers, enabling connection pooling and query caching for PostgreSQL and MySQL.

— Names supported databases and mechanism, no exact setup steps product docs

THINNER COVERAGE BELOW

05 **cloudflared drops 32-bit Windows and Intel macOS** DEPRECATED 55

Starting in 2027, `cloudflared` will stop publishing releases for 32-bit Windows and Intel-based macOS architectures.

— Gives exact timeline and affected architectures product docs

06 **RPC session spans in Workers traces** NEW 50

Automatically includes JavaScript RPC session spans in Workers traces for Worker-to-Worker and Worker-to-Durable Object calls.

— Automatic feature, no configuration step to act on product docs

07 **Inspect panel for Browser Run Session Recordings** NEW 45

Adds an Inspect panel to Browser Run Session Recordings, exposing console output, network activity, and page structure for post-session debugging.

— Names the panel but no navigation path or command product docs

08 **Reject-busy mode for synchronous inference** NEW 45

Adds a reject-busy mode for synchronous inference requests that fails immediately instead of waiting for capacity.

— Describes behavior but not how to enable it product docs

09 **Unified Routing GA in Cloudflare WAN** IMPROVED 30

Adds Unified Routing as generally available across supported WAN on-ramps in Cloudflare WAN.

— States GA status only, no mechanism or usage detail product docs

└──▷ **BREAKING ON UPGRADE**

! Starting in 2027, `cloudflared` will stop publishing releases for 32-bit Windows and Intel-based macOS architectures.

WAF - Enforce positive security with Application Profiles (2026-09-07) BLOG ↗

STABLE

Cloudflare WAF gains Application Profiles for positive-security enforcement, starting with Schema Profiles that learn and validate request structure.

└──▷ **HOW TO FIND IT**

Review Schema Profile conformance trends and violation samples in Security Analytics before committing to enforcement.

📍 In the Cloudflare dashboard, go to Security › Analytics › Profile Analysis to inspect conformance trends and sampled non-conforming requests for your Schema Profile.

Introduces Application Profiles, a new positive-security layer for Cloudflare WAF that learns expected request structure and classifies traffic as conforming or non-conforming without blocking.

Adds Schema Profiles as the first profile type, modeling path variables, query parameters, headers, cookies, JSON bodies, and form-encoded bodies — including field types, numeric ranges, string lengths, and character classes.

Adds Profile Analysis view in Security Analytics to review conformance trends and sampled violation details before enforcing a profile.

Enables enforcement of Application Profiles via Custom Rules, allowing scope by hostname, path, operation, or other security signals such as Attack Score.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 21, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue