THE CYBER TOOLCHAIN NO. 003

Head

THE WEEK'S TOP PICKS

PUBLISHED JULY 17, 2026 · FRIDAYS

// SHARE THIS WEEK ▽ Top 5 ▽ Top Pick

EDITIONS tail | grep | **head** | diff | uniq

The best tooling updates that shipped this week.

> // HOW THESE PICKS ARE MADE
>
> Every feature release from the 874 tools on our watchlist goes into the daily newsletter. Once a week we read the whole field side by side and pick the few updates worth your attention, on two questions: **how deep and complete is the single best capability in the release**, and **how much it changes what you can actually do**.
>
> A tool is judged on everything it shipped that week, so a project that releases daily gets credit for the sum — and still only takes one slot. We would rather run a short list than a padded one.

// THE PICKS

01	**Cloudflare**	Web Application Security
02	**Keycloak**	Identity & Access Management
03	**xsrfprobe**	Web Application Security
04	**Docker Sandboxes**	Runtime Security
05	**ssh-audit**	Vulnerability Management

01 Cloudflare

NEW TOOL

14 RELEASES · 2026-06-15 → 2026-07-14

Web Application Security

Cloudflare Blog is a publication sharing updates, technical insights, and announcements about Cloudflare's internet security, performance, and developer infrastructure products.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 5/5

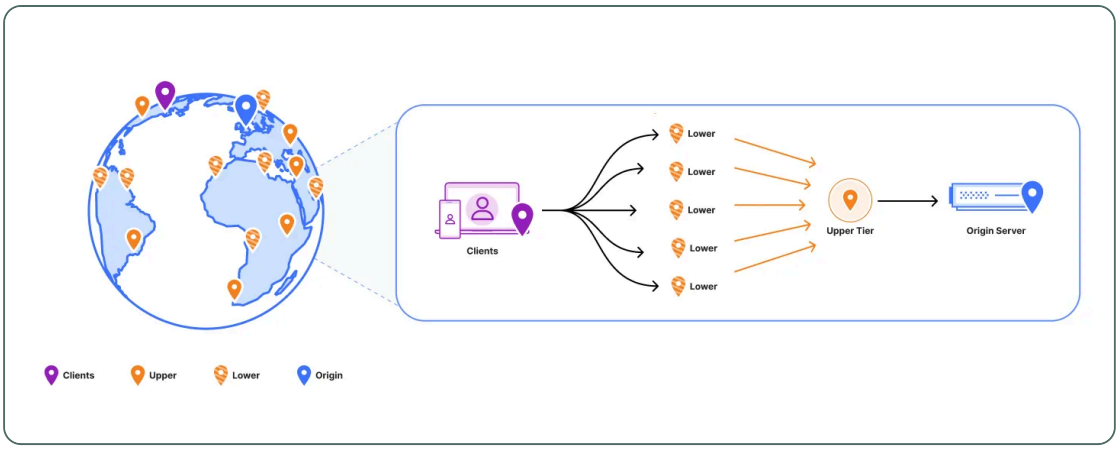
Precursor delivers session-scoped behavioral bot detection injected transparently into HTML responses — no app changes, continuous re-evaluation across the full user journey, and privacy-preserving signal collection — a capability that didn't exist in any WAF product before this week; combined with Cloudforce One threat-intel fields in WAF custom rules and the BotBase bot directory, this is the most complete and novel security release of the week.

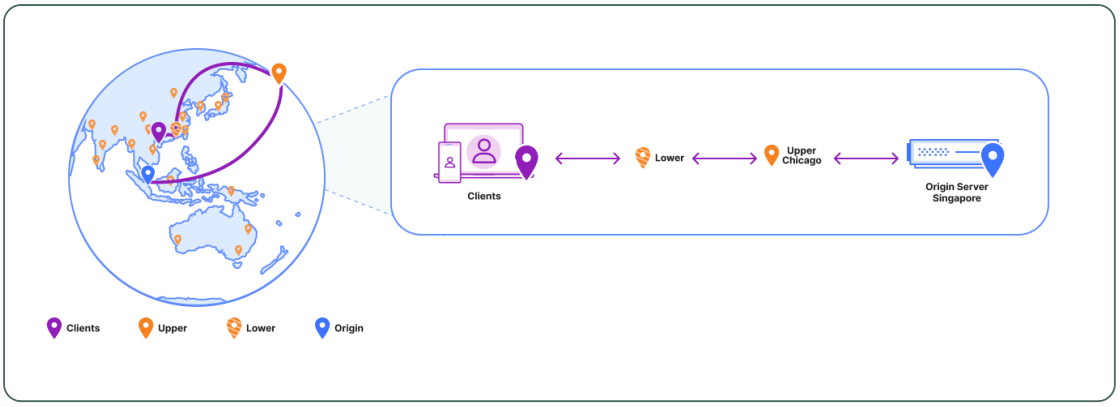
> Cloudflare shipped Precursor, a continuous client-side behavioral signal system that extends bot detection across entire user sessions, alongside a wave of WAF managed-rule detections (Citrix NetScaler, Kemp LoadMaster, FortiSandbox, Ivanti Sentry, Ghost CMS, and generic SSRF/LFI/XSS classes), new Cloudforce One threat-intelligence fields for WAF rules, EDE-based DNSSEC-bypass visibility in 1.1.1.1, cloud region hints for Smart Tiered Cache, and new AI-crawler traffic controls plus a BotBase bot directory and Attribution Business Insights dashboard.

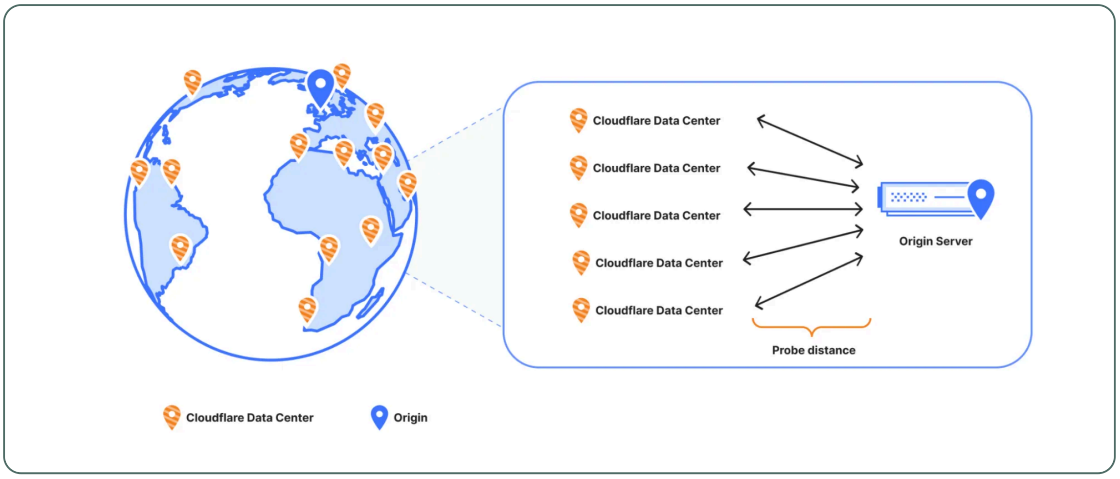
└──▷ WHAT SHIPPED · 12 FEATURES most completely described first ? what's the number?

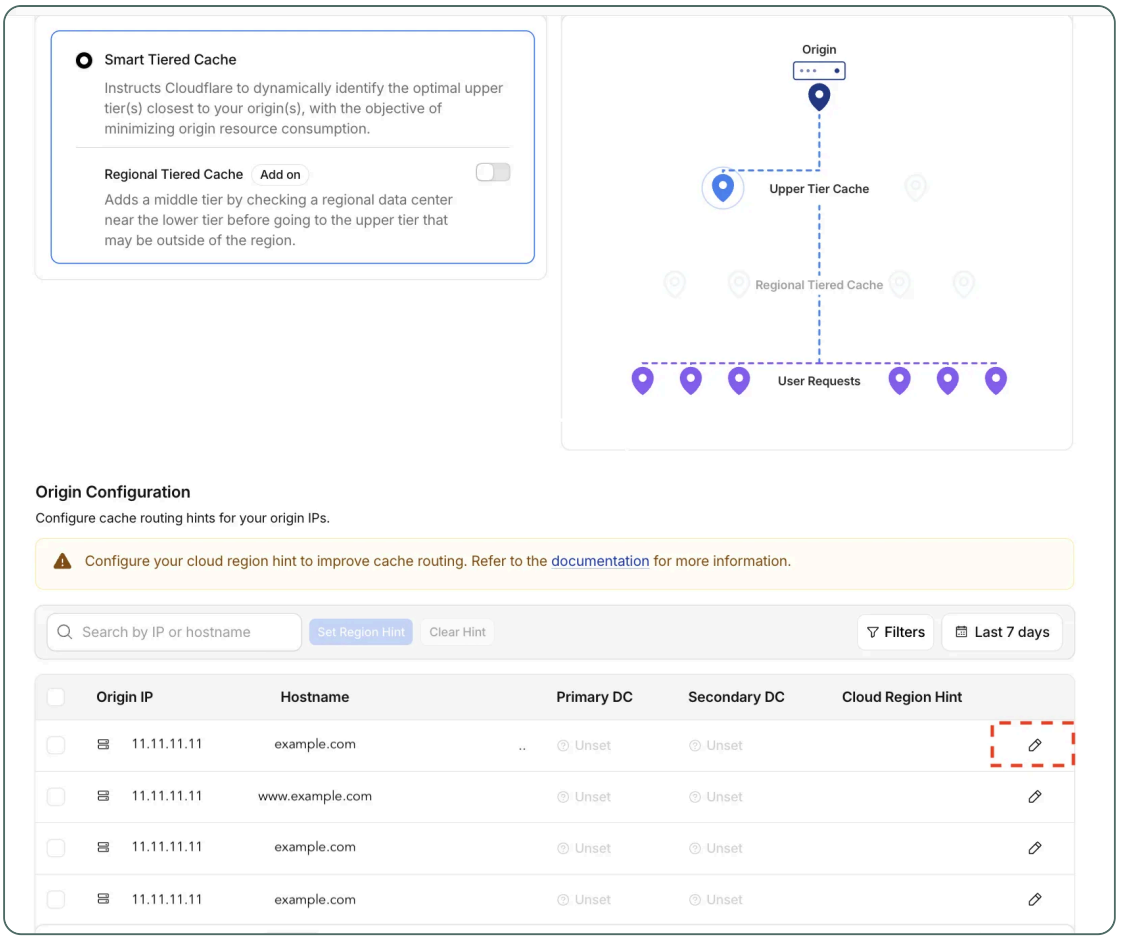
01 **Cloud region hints for Smart Tiered Cache** NEW 91

Smart Tiered Cache now supports cloud region hint configuration for anycast origins on AWS, GCP, Azure, and Oracle Cloud, letting it select the optimal upper tier instead of routing on default anycast geolocation. A new Origin Configuration table (dashboard: Caching > Tiered Cache > Origin Configuration) supports per-IP region hint assignment or bulk editing of all origin IPs, and region hints can also be managed via the Cloudflare API and Terraform. Anycast origins are detected using speed-of-light latency checks across multiple global checkpoint data centers, with cloud provider IP range mappings refreshed every few hours, upper-tier probe data refreshed every 15 minutes, and geographic fallback for regions lacking probe data.


Lower
Lower
Lower
Lower
Lower
Clients
Upper Tier
Origin Server
Clients
Upper
Lower
Origin


Clients
Lower
Upper
Chicago
Origin Server
Singapore
Clients
Upper
Lower
Origin


Cloudflare Data Center
Cloudflare Data Center
Cloudflare Data Center
Cloudflare Data Center
Cloudflare Data Center
Origin Server
Probe distance
Cloudflare Data Center
Origin



Set a region hint for an anycast AWS origin so Smart Tiered Cache routes through the optimal upper tier instead of hairpinning across continents.

> 📍 In the Cloudflare dashboard, go to Caching › Tiered Cache › Origin Configuration. Find your anycast origin IP, click the edit icon, select "Set Region Hint", and choose your region (e.g., aws:us-east-1). Save to apply.

— Exact UI path, API/Terraform support, and refresh intervals named.

Improving Smart Tiered Cache for Public …

02 Cloudforce One threat intelligence fields in WAF NEW

88

WAF custom and rate-limiting rules can now match new `cf.intel.ip.*` fields populated from Cloudforce One threat intelligence for IPs flagged within the past seven days: `cf.intel.ip.datasets` (DDoS or WAF attack activity), `cf.intel.ip.attacker_names` and `cf.intel.ip.attacker_countries` (known threat actor identity and source country), and `cf.intel.ip.target_industries` and `cf.intel.ip.target_countries` (historically targeted sectors/regions). Matches are recorded in Security Analytics and are fully supported via the Cloudflare API and Terraform.

Block inbound requests from IPs linked to DDoS campaigns that have targeted France — useful for protecting French-market infrastructure during active threat periods.

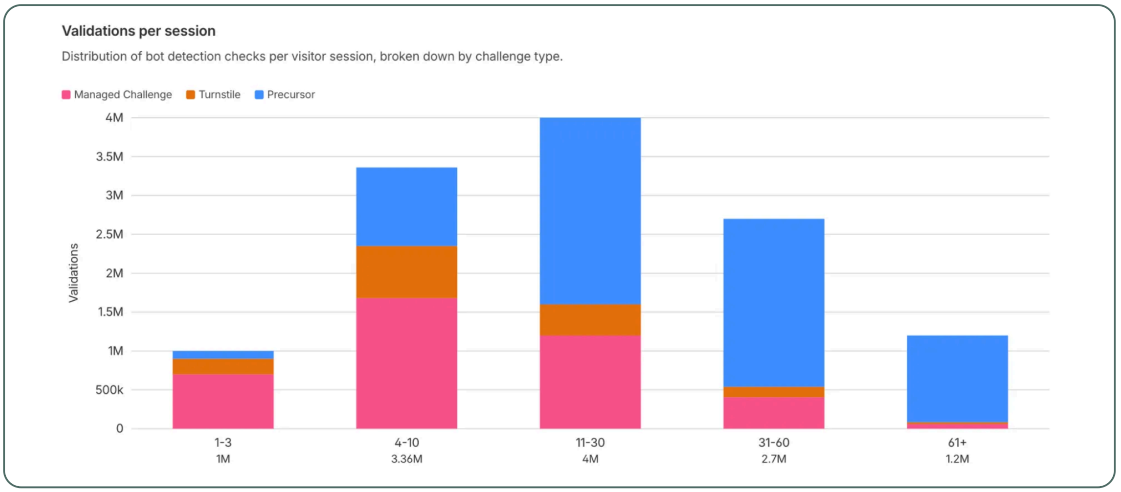
```
$ any(cf.intel.ip.target_countries[*] == "FR") and
any(cf.intel.ip.datasets[*] == "ddos")
```

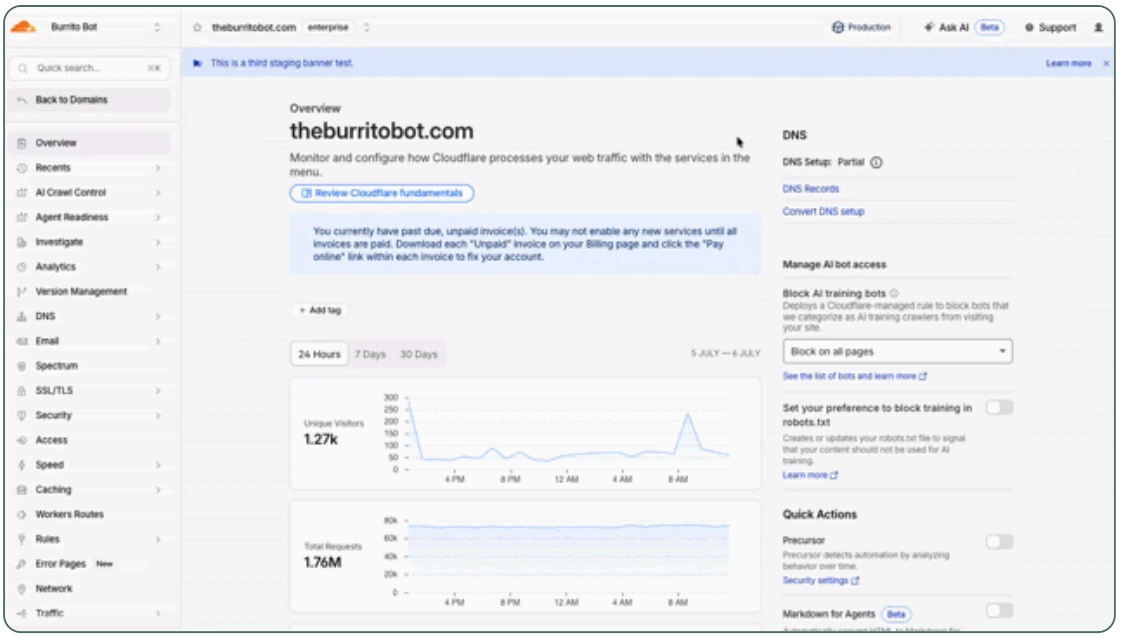
— Names every field and includes a runnable rule expression.

WAF - Use Cloudforce One threat intellig…

03 **Precursor session-based bot behavior detection** NEW 85

Precursor is a new client-side JavaScript engine Cloudflare dynamically injects into HTML responses to continuously collect behavioral signals — pointer movement, keyboard timing/rhythm, focus changes, visibility — across an entire user session rather than at a single challenge point, with no application changes required. Session-scoped evaluation accumulates signal throughout the session (preventing bots from resetting their signature via refresh or a new challenge) and cross-references multiple interaction streams in an edge-side evaluation layer with modular evaluators to detect automation such as linear mouse interpolation, inhuman click precision, missing physiological tremor, and atypical timing. Signals feed directly into existing bot score, challenge decisions, and Security Rules; a new session-based view in Security Analytics lets practitioners inspect full visitor journeys and find sessions with sustained automation. Keyboard data is captured only as timing/rhythm (not keystrokes) and is not exposed to dashboards or tied to identities. Precursor is enabled from the dashboard under Security > Bots, with a configurable mode (low-friction background observation or enforced Challenge for unverified sessions).





Enable Precursor on a zone to start collecting continuous session-level behavioral signals and feeding them into your existing bot defenses.

> 📍 In the Cloudflare dashboard, go to your zone › Security › Bots, find the Precursor section, toggle Precursor on, and select a verification mode (low-friction background observation or enforced Challenge for unverified sessions).

After enabling Precursor, investigate sessions showing signs of sustained automation — especially behavior occurring between individual requests.

> 📍 In the Cloudflare dashboard, go to Security › Security Analytics, then switch to the session-based view to filter bot management traffic by session-level signals and identify divergent visitor journeys.

Enable Precursor from the Cloudflare dashboard to activate session-based bot detection for your zone.

> 📍 In the Cloudflare dashboard, go to Security › Bots and enable Precursor, then select your preferred mode to balance security and user experience.

— Detailed mechanism and dashboard path, no API or flag named.

`Introducing Precursor: detecting agentic…` `Challenges - Precursor introduces sessio…`

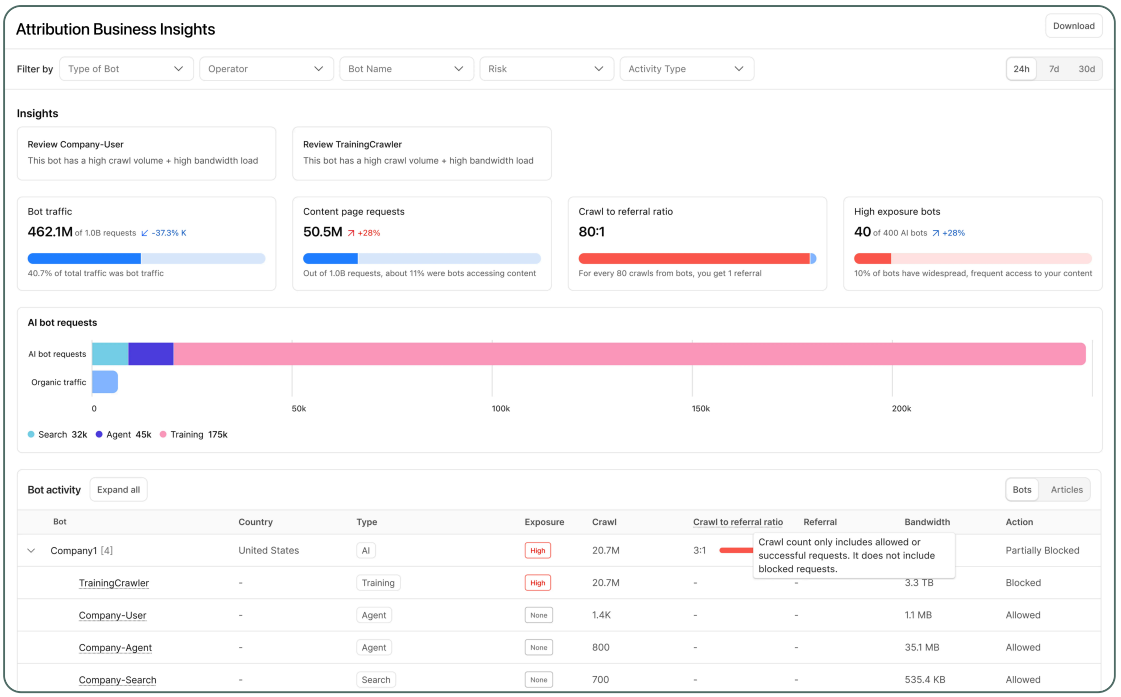
---

04 BotBase directory and Attribution Business Insights NEW

85

BotBase is a new searchable in-dashboard directory (Security > Bots > BotBase) of every bot Cloudflare tracks, classified by behavior (Search, Agent, Training, Transact, Data Collection, SEO, Ads Verification), letting practitioners filter live zone traffic to a single bot and copy its detection ID for use in Security > WAF > Custom Rules. Attribution Business Insights is a new dashboard showing crawl-to-referral ratios per bot operator

over 24 hours, 7 days, or 30 days with Allowed/Blocked/Partially blocked status, and every bot tracked in BotBase is also published to Cloudflare Radar's public bots and agents directory.



Use a bot's detection ID from BotBase to build a precise Security rule that blocks a specific crawler on your zone.

> 📍 In the Cloudflare dashboard, go to Security › Bots › BotBase, search for the target bot, copy its detection ID, then navigate to Security › WAF › Custom Rules and create a rule matching that detection ID.

Review which AI crawlers are sending referral traffic back versus just consuming your content, to inform allow/block decisions.

> 📍 In the Cloudflare dashboard, go to Security › Bots › Attribution Business Insights, set the time range to 30 days, and sort operators by crawl-to-referral ratio to identify high-crawl/low-referral operators.

— Exact dashboard paths and detection-ID workflow given.

`Bots - More visibility into bot traffic …`

---

05 **FortiSandbox path traversal detection** NEW 78

New managed rule `32075e19b1494117ac5915e8d84c92c9` blocks unauthenticated arbitrary file reads via Fortinet FortiSandbox path traversal (CVE-2026-39813), currently set to Log action.

Override the new FortiSandbox rule from Log to Block immediately if you run FortiSandbox and want active protection rather than waiting for Cloudflare's default escalation.

📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 32075e19b1494117ac5915e8d84c92c9, and set the action override to Block.

— Exact rule ID plus override-to-Block steps. `WAF - WAF Release - 2026-07-01 (2026-07-…`

06 **Ghost CMS and obfuscated boolean SQLi detections** NEW

78

New managed rule `439c4ef64b32447989bdf412b4c29bc6` detects and blocks blind SQL injection in the Ghost CMS Content API (CVE-2026-26980, versions 3.24.0–6.19.0), set to Block action. A second new rule `6c64b68ef5ed45e7a622cdaab56f403f` identifies generic SQLi bypass attempts using obfuscated boolean logic in URI parameters, shipped Disabled by default.

Enable the new obfuscated boolean SQLi rule (shipped as Disabled) to actively block evasion attempts against your zone.

📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 6c64b68ef5ed45e7a622cdaab56f403f (SQLi - Obfuscated Boolean - URI), and change its action from Disabled to Block.

— Both rule IDs named with exact version range and enable steps.

`WAF - WAF Release - 2026-06-15 (2026-06-…`

07 **Per-behavior AI crawler blocking controls** NEW

77

All plans, including Free, now get three independently controllable AI traffic behavior categories — Search, Agent, and Training — with per-behavior policies of block on all pages, block only on ad-displaying pages, or allow. A new default policy effective September 15, 2026 will automatically block Training and Agent crawlers on ad-displaying pages for newly onboarded domains (opt-out available before that date), and multi-purpose crawlers combining Search and Training behaviors will be subject to the Training-class blocking rules.

Opt out of the September 15, 2026 default changes before they apply to your newly onboarded domain.

📍 In the Cloudflare dashboard, go to Security › Bots › AI Traffic and manually configure each behavior preset (Search, Agent, Training) to your preferred setting before September 15, 2026.

— Names exact categories, policy options, and effective date.

`Bots - New options to manage AI traffic …`

08 **EDE codes for DNSSEC bypass in 1.1.1.1** NEW 75

1.1.1.1 now returns Extended DNS Error `EDE 33` (Negative Trust Anchor) on every response served while an NTA is active — including for non-DNSSEC domains — making DNSSEC validation bypass visible to clients and monitoring tools. It also returns `EDE 9` (DNSKEY Missing) alongside `EDE 33` to expose the underlying DNSSEC chain-of-trust failure that triggered the NTA.

Detect whether a response from 1.1.1.1 was served under a Negative Trust Anchor — useful in monitoring scripts or incident triage when a TLD-level DNSSEC rollover is suspected.

```
$ kdig @1.1.1.1 google.al
```

— Names exact EDE codes and includes a verification command.

A broken DNSSEC rollover took down .al. … | A broken DNSSEC rollover took down .AL. …

09 **Fake Bing/MSN Bot rule disabled by default** BREAKING 72

The managed rule `Anomaly:Header:User-Agent - Fake Bing or MSN Bot` (rule `ae20608d93b94e97988db1bbc12cf9c8`) has its action changed from Enabled/Block to Disabled; zones relying on it for bot enforcement must implement alternative coverage, such as a custom User-Agent matching rule.

Re-enable blocking of fake Bing/MSN bots via a custom WAF rule if your zone depended on rule ae20608d93b94e97988db1bbc12cf9c8 for bot enforcement.

> 📍 In the Cloudflare dashboard, go to Security › WAF › Custom Rules, create a new rule with field 'User Agent' matching known fake-bot patterns, and set action to Block.

— Exact rule ID and remediation steps for the breaking change.

WAF - WAF Release - 2026-07-01 (2026-07-…

10 **Citrix NetScaler and Kemp LoadMaster RCE detections** NEW 71

New Cloudflare Managed Ruleset detections block unauthenticated memory disclosure in Citrix NetScaler ADC and Gateway appliances configured as SAML IdPs (CVE-2026-8451, rule `78826e3223b94da493a2ade876973ac4`) and unauthenticated remote code execution via OS command injection in Progress Kemp LoadMaster load balancers (CVE-2026-8037, rule `6b64d216620449fbb273d07910233f36`), first scheduled for Log mode on 2026-07-13 and shipped live on 2026-07-14.

— Names both CVEs and rule IDs but no override example.

WAF - WAF Release - 2026-07-14 (2026-07-… | WAF - WAF Release - Scheduled changes fo…

11 **New SSRF, LFI, ColdFusion and XSS detections** NEW 70

Eight new managed detections deploy in Log mode on 2026-07-20: `SSRF - Restricted Protocol` (rule `7fbdc9407bdb4a4eae2b3d91215e7d31`), `SSRF - Obfuscated Host` (rule `6ca512d240d848d6a0c7ef42a935ee5d`), `LFI - Path Traversal` (rule `a3fb0870c38440d8a9a0eba81b0230ac`), Adobe ColdFusion File Upload Path Traversal CVE-2026-48276 (rule `452a04be3f73458c863d8dae61349c8b`), Adobe ColdFusion Path Traversal CVE-2026-48282 (rule `a53a3fb491c64d74908081ee9cb61eac`), and XSS via JavaScript bracket-concatenation obfuscation across three surfaces — request body (rule `d8b63828c2344d919b94d2594ac5e21f`), headers (rule `264a83a764be428ca41d516ff31f5559`), and URI (rule `4ba21a60837244029183b782987984fd`).

— Names every rule ID and CVE but no usage example.

WAF - WAF Release - Scheduled changes fo…

12 **Ivanti Sentry command injection detection** NEW 70

New Cloudflare Managed Ruleset detection (rule `500a90789f874345b60b0de7242fdf83`) blocks unauthenticated OS command injection attacks against Ivanti Sentry (CVE-2026-10520) at the WAF layer.

Verify the new Ivanti Sentry rule is active in Block mode on your zone — useful immediately after deployment to confirm coverage before patching.

> 📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 500a90789f874345b60b0de7242fdf83, and confirm the action is set to Block.

— Exact rule ID with verification steps included.

WAF - WAF Release - 2026-06-23 (2026-06-…

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

02 Keycloak

NEW FEATURE

3 RELEASES · 2026-07-03 → 2026-07-17

Identity & Access Management

Keycloak is an open-source identity and access management solution that provides single sign-on, authentication, and authorization for applications and services.

// WHY IT MADE THE LIST

DEPTH 5/5 IMPACT 4/5

Keycloak 26.7 ships stateless multi-cluster mode (auth sessions in DB, no external Infinispan required), a SCIM API preview, an SSF 1.0 Transmitter with CAEP/RISC event profiles and durable outbox, an AuthZen PDP endpoint, and SAML step-up auth — a release that advances the open-source identity stack on multiple fronts simultaneously.

> Keycloak's headline change is a database-backed stateless mode enabling Multi-Cluster v2 deployments without an external Infinispan cluster, accompanied by a wave of new experimental APIs — SCIM provisioning, a Shared Signals Framework transmitter, an AuthZen PDP, Identity Brokering v2, Admin API v2, ID-JAG, token exchange delegation, and organization-level fine-grained permissions.

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

01 Shared Signals Framework (SSF) Transmitter support NEW

96

Experimental SSF 1.0 Transmitter role pushes signed Security Event Tokens (SETs) to subscribed receivers for events like session revocation, credential change, and account disable, supporting CAEP 1.0 and RISC 1.0 event profiles and mapping native Keycloak events automatically. Delivery uses HTTP push (RFC 8935) and HTTP poll (RFC 8936) channels backed by a transactional outbox that persists events durably then drains asynchronously with exponential backoff; a synthetic event REST endpoint lets external tools inject events that didn't originate in Keycloak. Receivers are configured as standard OIDC clients with a per-client `ssf.enabled` toggle, `ssf.read` / `ssf.manage` scopes, a `defaultSubjects` policy of `ALL` or `NONE`, and a per-receiver `emit-only-events` allowlist; a legacy SSE CAEP draft profile supports Apple Business Manager/Apple School Manager interoperability. Full stream CRUD plus status, verification, and stream-config endpoints are exposed via the Admin REST API and new Admin Console SSF tabs, with Prometheus metrics under the `keycloak_ssf_*` namespace. Enable with `--features=ssf` or `--feature-ssf=enabled`.

Enable real-time SSF security event delivery so downstream apps learn about logouts and credential changes without waiting for token refresh.

```
$ bin/kc.sh start --features=ssf
```

Enable the experimental SSF Transmitter feature when starting Keycloak so the per-realm SSF endpoints become available.

```
$ kc.sh start-dev --feature-ssf=enabled
```

Obtain a client-credentials token scoped to SSF management, then use it to create or manage streams against the Keycloak SSF transmitter endpoints.

```
$ curl -s -X POST
"https://my.keycloak.test/realms/myrealm/protocol/openid-
connect/token" \
  -d "grant_type=client_credentials" \
  -d "client_id=caep-receiver" \
  -d "client_secret=…" \
  -d "scope=ssf.read ssf.manage"
```

Enable the SSF Transmitter for a realm and create an SSF-enabled OIDC receiver client, then configure its stream settings via the Admin Console.

> 📍 In the Admin Console: 1. Open the target realm → Settings → SSF Transmitter toggle → On (sets `ssf.transmitterEnabled`). 2. Go to Clients → Create client, set Client authentication On, Service accounts roles enabled. 3. On the client's SSF tab: enable SSF, set Default Subjects to ALL, add an Audience (e.g. https:/`/caep`.dev), and tick Push and Poll as supported delivery methods. 4. Under Client scopes, add ssf.read and ssf.manage as Optional scopes.

— Named endpoints, config keys, flags and full delivery/outbox mechanism

Keycloak 26.7.0 released (2026-07-09) | Experimental Shared Signals Framework su…

---

02 **Stateless mode and Multi-Cluster v2 deployment** NEW 91

New `stateless` preview feature stores authentication sessions, action tokens, and brute-force counters in the database instead of distributed Infinispan caches, enabling Multi-Cluster v2: two or more independent Keycloak clusters connected only through a synchronously replicated database, with no external Infinispan cluster required. Cross-cluster cache invalidation uses a database outbox/queuing table with a default polling interval of 100 ms; auth sessions survive full cluster restarts during minor upgrades. Failover is environment-agnostic, tested with Amazon Aurora PostgreSQL and AWS

Global Accelerator, removing the monitoring and failover/failback overhead of an external Infinispan cluster.



Start a single Keycloak node (or cluster) with stateless mode enabled to persist auth sessions and brute-force counters in PostgreSQL.

```
$ bin/kc.sh start \
  --features=stateless \
  --spi-cache-embedded--default--cluster-name=CLUSTER1 \
  --db=postgres --db-url=jdbc:postgresql://db-host/keycloak \
  --db-username=keycloak --db-password=<password>
```

Add a second independent cluster pointing at the same synchronously replicated database to achieve multi-cluster v2 failover with no Infinispan.

```
$ bin/kc.sh start \
  --features=stateless \
  --spi-cache-embedded--default--cluster-name=CLUSTER2 \
  --db=postgres --db-url=jdbc:postgresql://db-host/keycloak \
  --db-username=keycloak --db-password=<password>
```

— Full mechanism, tested stack, and runnable startup commands given

Multi-Cluster v2 and Stateless Mode now …

03 **OID4VCI credential refresh and issuance controls** IMPROVED

70

Adds OID4VCI credential refresh interval configuration via `vc.refresh_interval_in_seconds` in client scope settings, plus admin UI support for all OID4VCI settings and user-initiated credential issuance from the account console.

— Names exact config key and UI surfaces affected

Keycloak 26.7.0 released (2026-07-09)

04 **Identity Brokering API v2** NEW 62

Adds Identity Brokering API v2 with client-level authorization, confidential-clients-only enforcement, an OAuth 2.0-compliant POST endpoint, and a new session-based token storage option.

— Names enforcement model and new storage option, no exact path

Keycloak 26.7.0 released (2026-07-09)

05 **SCIM API for user/group provisioning** NEW 60

Preview SCIM API automates user and group provisioning and deprovisioning via standard CRUD, PATCH, filtering, pagination, and schema discovery endpoints; enabled via `--features=scim-api`.

Enable the SCIM API preview so HR systems or identity governance platforms can provision and deprovision users automatically.

```
$ bin/kc.sh start --features=scim-api
```

— Names endpoint operations and enabling flag but lacks exact paths

Keycloak 26.7.0 released (2026-07-09)

06 **AuthZen Evaluation API (PDP)** NEW 60

Experimental implementation of the OpenID AuthZEN Authorization API 1.0, exposing single and batch PDP evaluation endpoints for vendor-neutral authorization decisions; enabled via `--features=authzen`.

Enable the AuthZen PDP endpoint so PEPs can request standardized permit/deny decisions without coupling to Keycloak's proprietary authorization API.

```
$ bin/kc.sh start --features=authzen
```

— Named standard and flag but no endpoint paths given

Keycloak 26.7.0 released (2026-07-09)

THINNER COVERAGE BELOW

07 **Admin API v2 for declarative client management** NEW 57

Experimental Admin API v2 enables declarative OIDC and SAML client management with strict validation, an accurate OpenAPI spec, and Java, JavaScript, and CLI clients.

— Describes scope and clients but no config/flag given

Keycloak 26.7.0 released (2026-07-09)

08 **Token exchange delegation scope** NEW 48

Experimental token exchange delegation introduces a new delegation parameterized scope type that validates impersonation rights before granting token exchange.

— Names the scope type but no config or endpoint detail

Keycloak 26.7.0 released (2026-07-09)

09 **Identity Assertion JWT Authorization Grant (ID-JAG)** NEW 45

Experimental support allowing Keycloak to accept cross-domain identity assertions and issue access tokens without re-authentication.

— Explains purpose but no config surface or example given

Keycloak 26.7.0 released (2026-07-09)

10 **Step-up authentication for SAML clients** IMPROVED 45

Step-up authentication for SAML clients is promoted to supported status, enabling SAML service providers to request specific authentication context classes.

— Status change noted but no configuration steps given

Keycloak 26.7.0 released (2026-07-09)

11 **Fine-Grained Admin Permissions for organizations** NEW 43

Adds fine-grained delegation for organization administration, introducing dedicated admin roles and Fine-Grained Admin Permissions scoped to organizations.

— States capability without concrete role names or config

Keycloak 26.7.0 released (2026-07-09)

12 **MCP authorization documentation for Claude Code** NEW 25

Adds MCP authorization documentation for Claude Code as a new MCP client integration alongside existing VS Code support.

— Documentation-only mention with no technical detail

Keycloak 26.7.0 released (2026-07-09)

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

03 xsrfprobe

NEW FEATURE · 1 RELEASE · 2026-07-12

Web Application Security

The Prime Cross Site Request Forgery (CSRF) Audit and Exploitation Toolkit.

// WHY IT MADE THE LIST — DEPTH 4/5 — IMPACT 5/5

XSRFProbe v3 is a ground-up rebuild delivering 25 named, severity-rated CSRF checks covering token tampering, Origin/Referer bypasses, and SameSite bypass via client-side redirect gadgets and sibling-subdomain XSS — with optional headless Firefox for browser-dependent tests and auto-validated PoC generation; no other tool this week covers this attack surface at this depth.

> XSRFProbe v3.0.0 rebuilds its CSRF audit engine around 25 categorized checks with headless-browser SameSite bypass testing, a bounded crawler, JSON reporting, auto-validated proof-of-concept generation, and a response-diffing engine to cut false positives.

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

01 Headless browser SameSite validation and PoC auto-validation

NEW — 88

Adds optional headless Firefox (Selenium) integration via `--browser` to run browser-dependent SameSite bypass tests and auto-validate generated proof-of-concept pages; `--auto-validate-poc` triggers validation and `--geckodriver-path` points to a custom geckodriver binary. Each finding gets a tailored, auto-validated exploitable CSRF PoC page.

Run the full CSRF audit including browser-dependent SameSite bypass tests and get a machine-readable JSON report for pipeline consumption.

```
$ xsrfprobe --url https://target.example.com --browser --json
```

Auto-validate generated PoCs in headless Firefox and point to a non-PATH geckodriver binary for a custom test environment.

```
$ xsrfprobe --url https://target.example.com --browser --auto-validate-poc --geckodriver-path /opt/geckodriver/geckodriver
```

— Names flags and shows runnable commands using them v3.0.0

02 **Bounded crawler with scope limits** NEW 78

New crawler exposes configurable `--max-urls`, `--max-depth`, and `--crawl-timeout` flags to control scan scope and wall-clock time.

Limit crawler scope to a shallow pass on a large site so CI finishes quickly — cap URLs, depth, and wall-clock time.

```
$ xsrfprobe --url https://target.example.com --max-urls 50 --max-depth 2 --crawl-timeout 120 --json
```

— Exact flags with a runnable scoped-scan example v3.0.0

03 **JSON reporting for findings** NEW 78

Adds a `--json` flag producing machine-readable reports; each finding includes `test_id`, `severity`, and an `exploitability` note, suited for pipeline consumption.

Run the full CSRF audit including browser-dependent SameSite bypass tests and get a machine-readable JSON report for pipeline consumption.

```
$ xsrfprobe --url https://target.example.com --browser --json
```

— Named flag and fields with a runnable example v3.0.0

04 **SameSite bypass checks (S2–S4) and subdomain enumeration** NEW 75

New SameSite bypass checks `S2` – `S4` probe `SameSite=Strict` protections via client-side redirect gadgets and sibling-subdomain XSS, and `SameSite=Lax` via cookie-refresh/OAuth flows; sibling-subdomain testing is supported by optional crt.sh-based subdomain enumeration.

— Details bypass mechanisms and supporting enumeration, no command shown v3.0.0

05 **Token-tampering checks (T2–T8)** NEW 73

New token-tampering checks `T2` – `T8` actively probe GET↔POST method switching, token omission, empty token values, cross-session token replay, double-submit cookie bypasses, and custom-header token forgery.

— Lists exact bypass techniques and check IDs, no direct command v3.0.0

06 **Origin/Referer bypass checks (R0–R2c, O1–O3)** NEW 73

New Origin/Referer bypass checks `R0` – `R2c` and `O1` – `O3` test header removal, regex tricks, subdomain tricks, and `Origin: null` against real-world validation implementations.

— Names concrete bypass techniques and check IDs, no runnable command v3.0.0

THINNER COVERAGE BELOW

07 **25 categorized CSRF checks (D1–A1)** NEW 55

XSRFProbe's audit engine now runs a battery of 25 named checks spanning IDs `D1` – `A1`, covering token tampering, method/Content-Type overrides, Origin/Referer bypasses, SameSite cookie analysis, and token strength. Each check carries a unique `test_id`, a severity rating, and exploitability notes used in reporting.

— Names ID range and categories but no per-check mechanism v3.0.0

08 **Response diffing engine to reduce false positives** NEW 45

Adds a response diffing/benchmark engine that compares baseline and test responses before declaring an endpoint vulnerable, reducing false positives.

— Describes purpose but no mechanism detail or command v3.0.0

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

04 Docker Sandboxes NEW FEATURE

3 RELEASES · 2026-06-17 → 2026-07-10

Runtime Security

// WHY IT MADE THE LIST

DEPTH 4/5 IMPACT 4/5

Docker Sandboxes adds SOCKS5 egress proxy support, a native SSH endpoint for pairing-free access, kit source allowlisting, and a secret import workflow that keeps credentials out of the sandbox environment — and closes a DNS-based data-exfiltration CVE (CVE-2026-12039) by gating DNS resolution on network policy — a runtime security sandbox that meaningfully advances the practitioner's ability to contain agentic workloads.

> Docker Sandboxes shipped three releases (v0.33.0-v0.35.0) that add a SOCKS5 proxy transport for sandbox egress, an experimental native SSH endpoint, kit source allowlisting, wildcard host patterns for secrets, named sandbox identity/re-attachment, and network-policy hardening that closes two CVEs, alongside a long tail of policy, kit, and CLI improvements.

└──▷ WHAT SHIPPED · 28 FEATURES most completely described first ? what's the number?

01 Native SSH endpoint for sandboxes NEW

88

A new experimental native SSH endpoint in sandboxd allows connecting with `ssh <sandbox-name>@127.0.0.1 -p 2222` (publickey auth, connect-to-create, interactive shell and exec), enabled via `sbx settings set feature.ssh true`.

Enable the experimental SSH endpoint so agents or operators can shell directly into a running sandbox without a separate proxy.

```
$ sbx settings set feature.ssh true
```

Connect an interactive shell to a named sandbox over the native SSH endpoint once feature.ssh is enabled.

```
$ ssh <sandbox-name>@127.0.0.1 -p 2222
```

— Exact enable and connect commands provided. v0.34.0

02 SOCKS5 proxy support for sandbox egress NEW

86

Adds a SOCKS5 upstream-proxy transport (`socks5://` / `socks5h://`) for sandbox egress, configurable via `DOCKER_SANDBOXES_PROXY` , `HTTP_PROXY` , or `HTTPS_PROXY` , with optional proxy authentication and exclusions via `DOCKER_SANDBOXES_NO_PROXY` .

Route all sandbox egress through a corporate SOCKS5 proxy (with credentials) so sandbox traffic flows through your org's inspection infrastructure.

```
$
DOCKER_SANDBOXES_PROXY=socks5://proxyuser:proxypass@proxy.corp.e
xample.com:1080 DOCKER_SANDBOXES_NO_PROXY=localhost,127.0.0.1
sbx run
```

— Full env vars and a runnable command given. v0.35.0

---

03 **Kit source allowlisting** BREAKING 86

Kit installs are now restricted to an allowlist of sources, defaulting to Docker Hub (`docker.io/`) only; installing from another registry or a Git URL fails until its prefix is added with `sbx settings set kit.allowedSources` .

Allow kit installs from a private registry or Git source in addition to Docker Hub, to avoid the new allowlist breaking change.

```
$ sbx settings set kit.allowedSources ghcr.io/myorg/
```

— Exact config command and default scope given. v0.34.0

---

04 **Wildcard host patterns for secrets** NEW 86

`sbx secret set-custom --host` and `serviceDomains` in kits now accept wildcard host patterns (`*` for one label, `**` for any number of labels) and are repeatable, allowing one secret to cover multiple subdomains.

Grant a custom secret to all subdomains of an internal service using wildcard host patterns.

```
$ sbx secret set-custom --host '*.internal.corp' --host
'**.api.internal.corp' MY_API_TOKEN
```

— Exact syntax and runnable example command given. v0.33.0

---

05 **Policy CLI: inspect, check, filtered listing, and init rename** IMPROVED 85

New `sbx policy inspect` examines policy details, `sbx policy check network` tests whether the current policy would allow a network access request before executing, and `sbx policy ls` gains `--wide`, `--source`, and `--decision` filters. `sbx policy set-default` is renamed to `sbx policy init` (old name kept as a hidden deprecated alias).

Before running an agent task, verify that the current policy permits a specific outbound connection — catch misconfigurations without launching the full sandbox.

```
$ sbx policy check network --host api.example.com --port 443
```

— Names every subcommand and flag, with a runnable example. v0.35.0 v0.34.0

06 **Secret import and keychain migration** BREAKING 83

New `sbx secret import` imports credential env vars into the keychain, and `sbx secret ls` now flags env-only and OAuth-shadowed entries. Host env vars no longer auto-inject at runtime, requiring migration via `sbx secret import`.

Migrate host env-var credentials (e.g. API keys) into the keychain so they are explicitly managed rather than auto-injected at runtime.

```
$ sbx secret import
```

— Command given plus clear breaking-change migration path. v0.35.0

07 **Named sandbox identity and re-attachment** IMPROVED 80

A new `--name` flag for `sbx run` identifies a sandbox independently of the working directory, enabling multiple named sandboxes in the same workspace and re-attachment from any directory. Re-attaching to a sandbox created with a custom `--kit` agent now works via `sbx run --name <name>` without re-passing `--kit`.

Run two independent Claude Code sandboxes in the same workspace — useful when parallelising agent tasks without them colliding.

```
$ sbx run --name feature-a claude && sbx run --name feature-b
claude
```

Re-attach to a named sandbox from a different directory without specifying an agent, useful when returning to a long-running session.

```
$ sbx run --name feature-a
```

— Runnable commands shown for creating and re-attaching. v0.33.0 v0.35.0

08 **Network policy hardening: DNS gating and ICMP blocking** IMPROVED 68

Sandbox DNS resolution is now gated on network policy, closing a DNS-based data-exfiltration channel (CVE-2026-12039), with loopback names remaining exempt. Outgoing ICMP from sandboxes is now blocked persistently across daemon restarts (CVE-2026-12539).

— Two CVEs named with scope; no user-facing config. v0.33.0

09 **virtiofs caching enabled by default** IMPROVED 60

virtiofs cache is now enabled by default on macOS and Linux for improved filesystem performance, and this was extended to all operating systems; opt out via `DOCKER_SANDBOXES_ENABLE_VIRTIOFS_CACHE=0`.

— Opt-out env var named; no performance numbers given. v0.35.0 v0.33.0

THINNER COVERAGE BELOW

10 **`sbx kit add` recreates sandbox with full kit set** IMPROVED 58

`sbx kit add` now recreates the sandbox container with the full kit set while preserving state, and applies the kit's network allow/deny rules and composed policy.

— Explains before/after behaviour but no example. v0.35.0

11 **Daemon lifecycle management CLI** NEW 55

New `sbx daemon` command with `start`, `stop`, `status`, and `log-level` subcommands for managing the sandbox daemon's lifecycle.

— Names subcommands but no usage example given. v0.35.0

12 **OCI v2 kit artifact streaming** IMPROVED 55

Adds OCI v2 kit artifact streaming that decompresses layers once to a cache directory and uses seek-based random access, reducing memory pressure during kit reads.

— Mechanism explained but nothing for reader to invoke. v0.34.0

13 **`sbx rm --force` required for active sessions** BREAKING 50

`sbx rm` now requires `--force` to delete an active session, preventing accidental teardown.

– Exact flag named, no example command shown. v0.35.0

14 **Balanced network preset domain additions** IMPROVED 50

The balanced network preset now allows VS Code domains, Azure Blob Storage (`*.blob.core.windows.net`), and `dhi.io` over HTTP.

– Named domains but no config steps. v0.35.0

15 **`sbx ls --json` stable sandbox `id` field** IMPROVED 50

`sbx ls --json` now reports a stable per-sandbox `id` field for scripting and automation.

– Exact flag and field named, no example output. v0.33.0

16 **Claude connectors work without policy overrides** IMPROVED 50

Claude connectors (Slack, Gmail, Notion, Atlassian, etc.) now work inside sandboxed Claude Code without manual policy overrides.

– Names specific connectors but no config steps. v0.33.0

17 **Keychain fallback to encrypted on-disk store** NEW 50

New keychain fallback to an encrypted on-disk store on Linux when `dbus-launch` is unavailable, enabling headless/server deployments.

– Trigger condition named, no setup steps. v0.33.0

18 **MITM proxy CRL support** NEW 50

MITM proxy now publishes a CRL and embeds a CRL distribution point in generated certificates, supporting revocation-strict TLS clients (e.g., .NET).

– Mechanism named, no configuration for reader. v0.33.0

19 **`sbx inspect` surfaces kits and secrets** IMPROVED 45

`sbx inspect` now surfaces the sandbox's kits, injected secrets, and sandbox information.

– States what's shown, no example or output format. v0.35.0

20 **Docker login token injection via `sbx-login` credential** NEW 45

Kits can inject the user's Docker login token into requests to docker.com hosts via a credential with service `sbx-login`.

— Names the credential service but no usage steps. v0.35.0

21 **Faster sandbox startup via TLS-proxy CA optimization** IMPROVED 45

Sandbox startup is faster due to optimized TLS-proxy CA installation, which merges into the trust bundle instead of running `update-ca-certificates`, saving several hundred milliseconds.

— Mechanism and rough timing given, no user action. v0.35.0

22 **`sbx setup` onboarding command** NEW 45

New experimental `sbx setup` command imports agent credentials from environment variables for smoother first-time onboarding.

— Command named, no example invocation shown. v0.34.0

23 **`sbx cp -L` follows symlinks** IMPROVED 45

`sbx cp -L` now follows symlinks in the source path for sandbox-to-host copies.

— Exact flag named, minimal further detail. v0.33.0

24 **Sandbox port restoration and conflict recovery** IMPROVED 35

Restores published sandbox ports on restart, with automatic recovery from host-port conflicts by selecting a new host port.

— Describes behaviour, no command or config named. v0.34.0

25 **Droid OAuth credentials proxy-managed** IMPROVED 30

Droid OAuth credentials are now proxy-managed, keeping real tokens on the host and out of the sandbox.

— States benefit but no mechanism or config. v0.35.0

26 **Linux arm64 build packages** NEW 30

Linux `arm64` build packages are now produced, extending platform support.

— States new platform, no install instructions. v0.33.0

27 **`sbx exec` matches `sbx run` working directory** IMPROVED 30

`sbx exec` now uses the same working directory as `sbx run`.

— One-line consistency fix, no further detail. v0.33.0

28 **Cursor OAuth support** NEW 25

Adds Cursor OAuth support as a recognized agent.

— Bare statement, no configuration detail. v0.33.0

└──▷ BREAKING ON UPGRADE

! Host env vars no longer auto-inject at runtime — use `sbx secret import` to migrate existing env-var-based secrets.

! `sbx rm` now won't delete an active session unless `--force` is passed.

! Installation is blocked on Windows versions older than Windows 11.

! Kit installs are restricted to an allowlist of sources defaulting to Docker Hub (`docker.io/`) only; installing a kit from another registry or a Git URL fails until you add its prefix with `sbx settings set kit.allowedSources` .

! The `sbx policy set-default` command is renamed to `sbx policy init` (the old name still works as a hidden deprecated alias, but is deprecated).

! The `sbx run <name>` positional argument form for re-attaching to an existing sandbox now prints a deprecation warning; the positional argument for `sbx run` is expected to be an agent (e.g. `claude` , `codex`), not a sandbox name, and the old positional-name form will be removed in a future release.

! `sbx run --name` no longer auto-creates numbered sibling sandboxes when a same-named sandbox exists.

! `sbx create` now errors when run without arguments instead of proceeding.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

05 ssh-audit

NEW FEATURE · 1 RELEASE · 2026-07-04

Vulnerability Management

SSH server & client security auditing (banner, key exchange, encryption, mac, compression, compatibility, security, etc)

// WHY IT MADE THE LIST — DEPTH 3/5 — IMPACT 4/5

ssh-audit v3.9.0 adds post-quantum gap warnings (flagging every key exchange lacking PQC protection), UNIX socket scanning for containerized environments, SOCKS5 proxy support, and built-in offline hardening guides — making it immediately actionable for practitioners auditing SSH posture ahead of quantum migration deadlines.

> ssh-audit v3.9.0 adds post-quantum key exchange warnings, UNIX socket and SOCKS5 proxy scanning, and built-in offline hardening guides, alongside expanded algorithm detection and the removal of SSHv1 and older Python support.

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

01 Built-in offline hardening guides NEW

80

New `--list-hardening-guides` and `--get-hardening-guide` flags surface built-in hardening guides for common platforms (Debian, Ubuntu, Rocky Linux, OpenSSH) without requiring internet access.

Retrieve the built-in hardening guide for a specific platform to apply recommended SSH configuration without visiting an external site.

```
$ ssh-audit --list-hardening-guides
ssh-audit --get-hardening-guide "Debian 13"
```

— Exact flags and runnable commands for retrieving guides offline. v3.9.0

02 SOCKS5 proxy scanning NEW

80

New `--socks5 host:port` flag enables scanning through a SOCKS5 proxy; using it implies `--skip-rate-test`.

Scan an internal host through a SOCKS5 jump proxy — useful when direct connectivity to the target is unavailable.

```
$ ssh-audit --socks5 proxy.example.com:1080 target-host.internal
```

— Exact flag, implied behavior, and runnable command given. v3.9.0

03 **UNIX socket scanning** NEW 70

ssh-audit supports scanning SSH servers over UNIX sockets (e.g. `unix:///path/to/socket`), enabling auditing in containerized and CI/CD environments.

Audit an SSH daemon exposed only via a UNIX socket inside a container or CI/CD pipeline — no TCP port required.

```
$ ssh-audit unix:///run/ssh-unix-local/socket
```

— Named socket path format with a runnable example command. v3.9.0

04 **SSHv1, Python 3.8/3.9, and batch-mode verbosity removed** BREAKING 70

SSHv1 support has been removed entirely, breaking scans or configurations that relied on it. Support for Python 3.8 and Python 3.9 has been dropped. Batch mode (`-b`) no longer implicitly enables verbose mode, so scripts depending on that behavior must now add `-v` explicitly.

— Names exact flags and versions dropped, with clear migration note. v3.9.0

05 **Expanded algorithm identification** NEW 60

ssh-audit now identifies 2 new post-quantum key exchanges (`mlkem768nistp256-sha256` , `mlkem1024nistp384-sha384`), 11 new host key types including ML-DSA variants (`mldsa-44` , `mldsa-65` , `mldsa-87` , and OpenSSH-namespaced equivalents) and `webauthn-sk-ecdsa-sha2-nistp256-cert-v01@openssh.com` , and 2 new ciphers (`AEAD_CAMELLIA_128_GCM` , `AEAD_CAMELLIA_256_GCM`).

— Names every new algorithm but no way to act beyond detection. v3.9.0

06 **Policy and platform guide updates** IMPROVED 60

Adds an `allow_hostkey_subset_and_reordering` policy option letting policy definitions accept targets with a stricter or differently-ordered host key list, plus new built-in policies for OpenSSH 10.0 through 10.4 and new platform hardening guides for Debian 13, Rocky Linux 10, and Ubuntu 26.04.

— Names the config key and covered versions but no usage command. v3.9.0

THINNER COVERAGE BELOW

07 **Post-quantum key exchange warnings** NEW 50

ssh-audit now warns on all key exchanges that lack post-quantum protection, flagging exposure to Harvest Now, Decrypt Later attacks.

— Explains the risk flagged but no flag or output format given. `v3.9.0`

└──▷ BREAKING ON UPGRADE

! Support for Python 3.8 has been dropped, so setups running ssh-audit on Python 3.8 will no longer work.

! Support for Python 3.9 has been dropped, so setups running ssh-audit on Python 3.9 will no longer work.

! SSHv1 support has been removed, so any scans or configurations relying on SSHv1 scanning will no longer function.

! Batch mode (`-b`) no longer automatically enables verbose mode, so scripts or workflows that depended on batch mode implying verbose output will see less output unless `-v` is explicitly added.

GOOD PICK? 👍 👎 GOOD REASON? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · This Week's Highlights · July 17, 2026
The full firehose — every feature release, every day — is in the newsletter.

▷ Subscribe