THE CYBER TOOLCHAIN

NO. 062

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 2, 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 190 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 40 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 | Opus 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 today: Elastic Security Labs can turn Sentinel migrations and alert triage into assisted SOC work, cve-lite-cli

vim ? help

ranks CVEs by exploit likelihood instead of severity alone, and JFrog ties signed release evidence to compliance frameworks.

RESPOND

Move alert triage and detection migration out of copy-paste mode

Elastic Security 9.5 adds AI-assisted Sentinel rule migration, alert triage, and Attack Discovery workflows that can draft detection rules. That reduces the analyst toil around porting detections and turning an investigation into reusable coverage.

Elastic Security Labs

DETECT

Prioritize vulnerabilities by exploitation likelihood, not severity alone

cve-lite-cli adds an EPSS Priority Signal that blends CVSS with FIRST.org exploitation-likelihood data into four tiers, while Checkmarx adds native reachability analysis inside SCA scans. Teams get a cheaper first cut on which findings are likely exploitable or reachable before spending remediation time.

cve-lite-cli · Checkmarx

GOVERN

Attach evidence and policy to releases and AI components

JFrog now auto-assembles signed release-evidence chains for NIST SSDF and EU CRA reporting, and expands governance over MCP servers, skills, plugins, and agent packages. Aembit adds policy inspection for MCP traffic enriched with workload identity, giving teams a control point for AI-agent activity instead of treating it as opaque application traffic.

JFrog · Aembit

DETECT

Find more attack surface without waiting for cloud access

Noir extends endpoint detection into CLI frameworks across 20+ languages and can run passive-scan rules offline with custom rule paths. Prowler adds 25 checks across AWS and seven other providers plus Slack alerting, making newly discovered cloud and API findings easier to surface without manual report polling.

noir

Does **GPT-5.5** read better? 👍 👎

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-09-01

NOTES ↗

width=device-width,initial-scale=1" name="viewport"/> <meta content=

> NodeZero's 2026.08 release adds two new pentest op types — WebApp and Azure Cloud — alongside Australian data residency, new Activity & Artifacts tabs for distinguishing NodeZero actions from real attacks, expanded Rapid Response coverage, and several Active Directory & Identity assessment upgrades.

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

01 **Azure Cloud pentest op type** NEW 60

A new Azure Cloud pentest op type combines Azure API enumeration with an Internal pentest for a single assumed-breach assessment covering both Azure control planes and workloads running on Azure VMs.

Assess an Azure environment end-to-end — from control plane enumeration to lateral movement on VMs — in a single assumed-breach op.

> 📍 In NodeZero, go to Launch Op › Azure Cloud Pentest, supply your Azure credentials/scope, and select the combined Azure API + Internal pentest assessment mode.

— Clear mechanism and scope; only UI launch steps as action, no API. snapshot-20260901

02 **Activity & Artifacts tabs across Real-Time and post-op views**

NEW 60

New Activity and Artifacts tabs are added to both the Real-Time View and post-op test results; the Action Log is extended across both views, and the Artifacts tab lets reviewers browse collected files to correlate NodeZero's actions against IDS/SIEM alerts and rule out false positives, helping distinguish NodeZero activity from malicious attacks.

After an op completes, use the Artifacts tab to review files NodeZero collected and correlate its actions against IDS/SIEM alerts to rule out false positives.

> 📍 In the NodeZero console, open a completed pentest, select the Activity tab, then switch to the Artifacts tab to browse collected artifacts alongside the Action Log entries.

— Exact UI navigation given; names the tabs and their purpose. snapshot-20260901

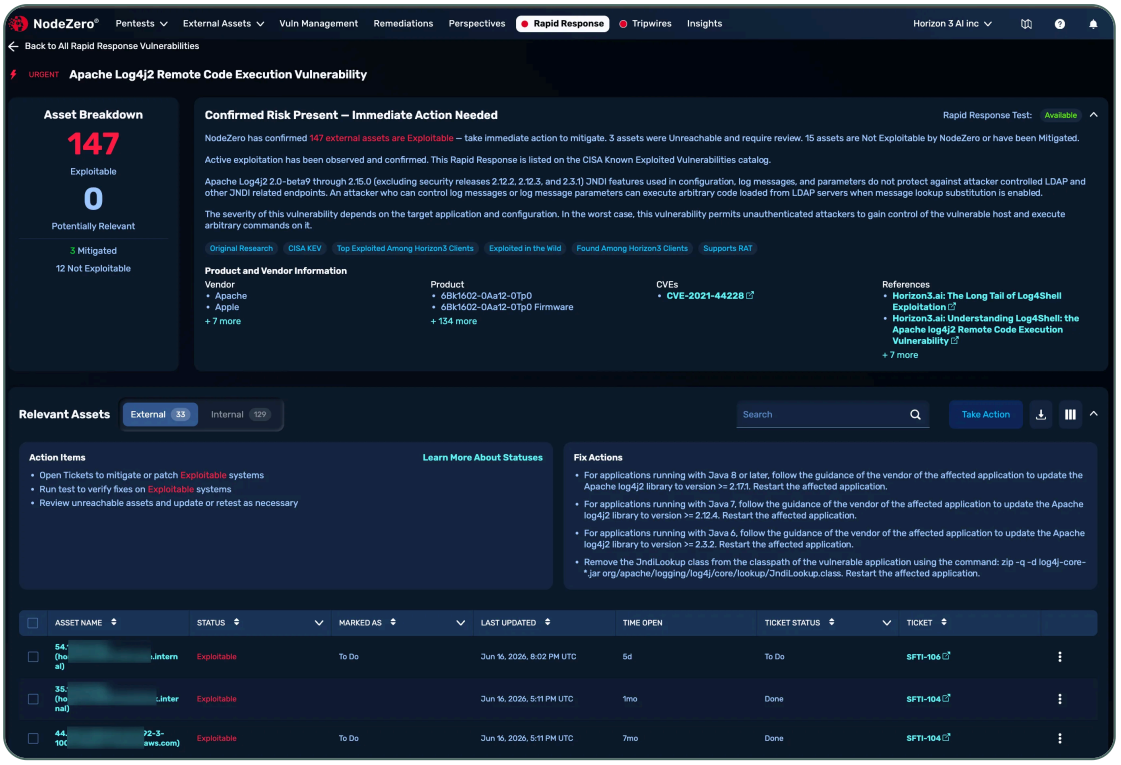
THINNER COVERAGE BELOW

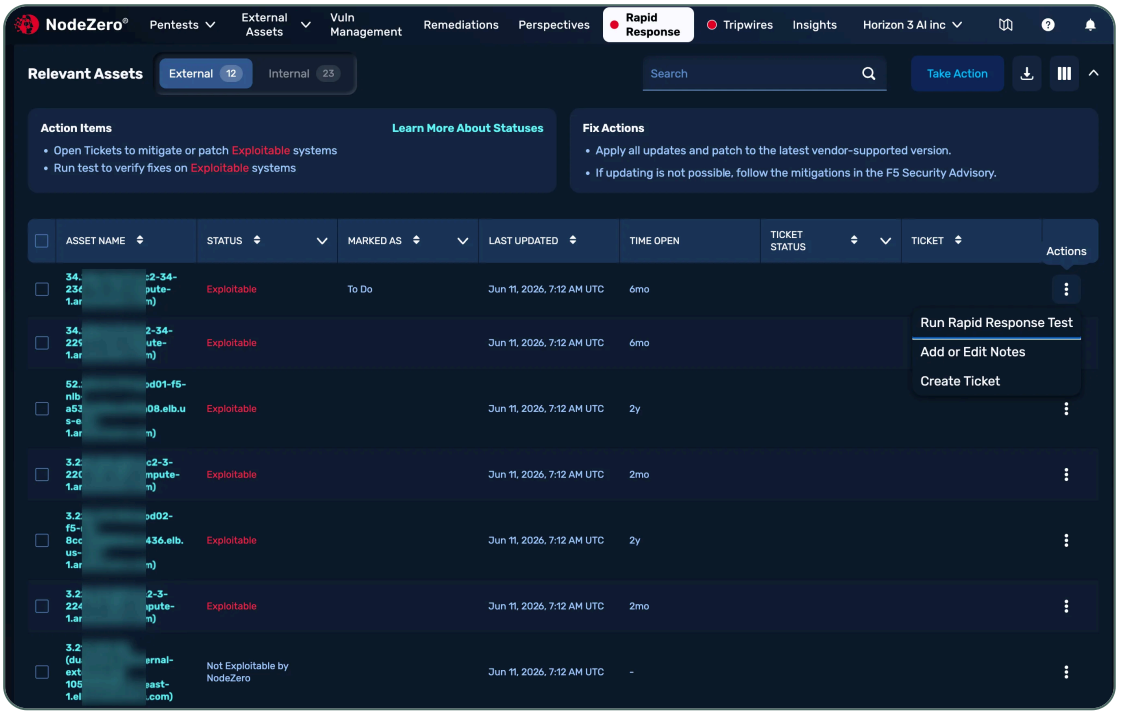
03 Rapid Response tests for Metabase, GitLab, Ivanti, plus cert and S3 takeover detection

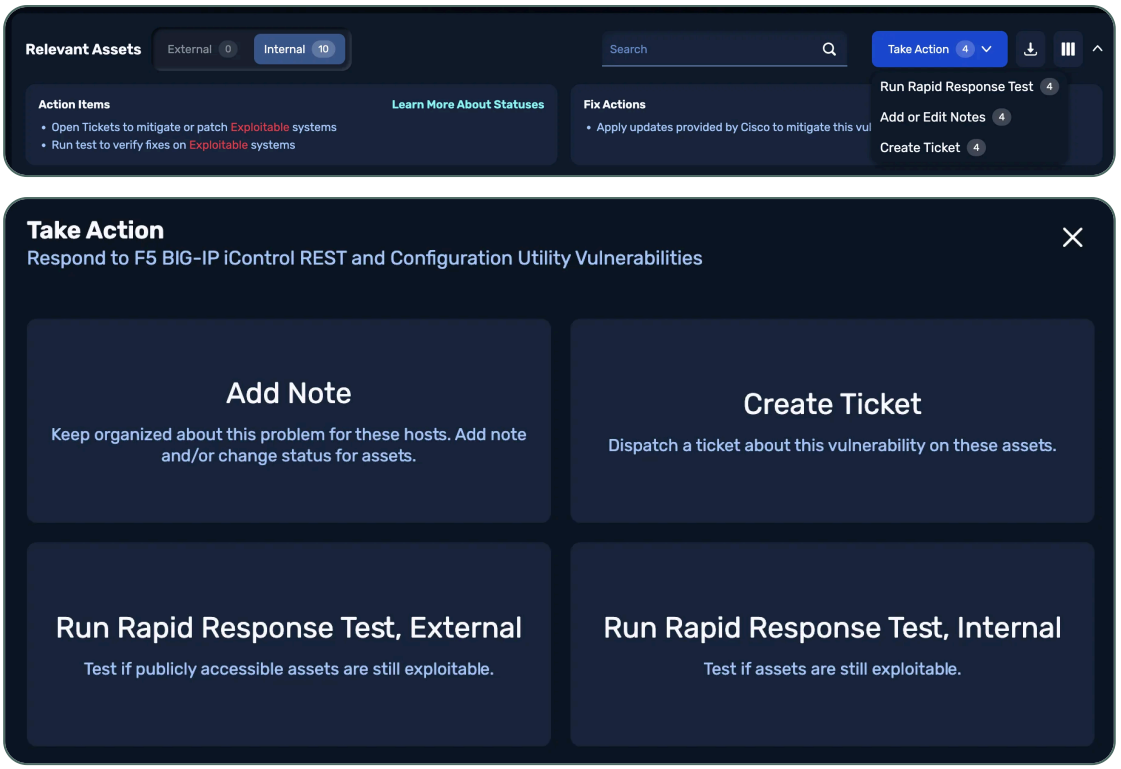
NEW

55

New targeted Rapid Response tests cover critical vulnerabilities in Metabase, GitLab Community and Enterprise Editions, and Ivanti Endpoint Manager, plus expanded certificate and Amazon S3 takeover detection.







— Names all affected products but gives no steps to run a test. `product docs`

---

04 **Active Directory & Identity assessment upgrades** IMPROVED

50

Active Directory and Identity assessments gain a dedicated Identity Details view, Kerberos ticket extraction during credential dumping, and safer authentication workflows.

— Names three concrete additions but no UI path or command given. `snapshot-20260901`

---

05 **Australia data residency** NEW

30

Platform traffic and pentest data can now be kept within the Australian region for data-sovereignty compliance.

— States the capability but no setup steps or config named. `snapshot-20260901`

WAS THIS USEFUL?

adscan

Feedback

1 RELEASE · 2026-09-01

NOTES ↗

Free Active Directory pentesting tool for Linux. Automates AD and LDAP enumeration, Kerberoasting, ASREPRoast, password spraying, ADCS/ESC1, DCSync, RBCD, gMSA, shadow credentials, NTLM relay and credential dumping across 104 techniques, mapping BloodHound-compatible attack paths to Domain Admin.

> ADscan v11.3.0 adds a didactic teaching mode, seasonal password spraying, a minimal-footprint krbtgt-only DCSync option, cross-forest attack path analysis with two automated escalation routes, LDAP credential-attribute harvesting as a full attack step, and pre-filled independent-verification commands for every finding.

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

01 **Independent verification commands per finding** NEW

85

Reports and the web platform now include a Verify this finding independently block per attack step, with native Windows/PowerShell (Get-ADUser , Get-Acl , `dsacls` , `certutil`) and Linux equivalents pre-filled with the real DC IP and domain name, covering Kerberoasting, AS-REP Roasting, DCSync, ADCS ESC1, and the GenericAll/GenericWrite/WriteDACL/WriteOwner ACL family.

— Named exact commands and coverage, pre-filled and ready to run v11.3.0

02 **Didactic mode for technique explanations** NEW

80

Adds `set explain_level off|basic|deep` to control didactic mode verbosity, and `explain <technique>` (e.g. `explain kerberoasting` , `explain "ADCS ESC1"`) to display a teaching card — MITRE ATT&CK ID, equivalent manual command, and detection Event IDs — for any technique without executing it.

Before a Kerberoasting step in a lab, display its full teaching card — MITRE ID, manual equivalent, and detection Event IDs — without triggering the attack.

```
$ explain kerberoasting
```

— Names exact commands and card contents; runnable as shown v11.3.0

03 **Seasonal password spraying strategy** NEW

80

Adds a `month_season` spray strategy to the scan config for seasonal password spraying (e.g. March2026, Verano2026), with per-domain language inference (English/Spanish) and last-password-change awareness when domain credentials are available; runs automatically in `adscan ci` .

Run a fully unattended CI scan that includes both month and season spray passes against a domain, using the month_season strategy.

```
$ adscan ci auth --type audit --interface eth0 --domain
corp.local --dc-ip 10.0.0.1 -u alice -p 'S3cr3t!'
```

— Named config strategy and CI integration with a runnable example v11.3.0

04 **Automated cross-forest escalation routes** NEW 65

Automated cross-forest escalation now ships two routes: from a compromised trusted forest into the trusting forest via cross-organization TGT-delegation trust, and from a compromised child domain up to its forest root (RaiseChild); each route triggers post-compromise actions on the reached domain.

— Clear before/after mechanism, no exact command shown v11.3.0

05 **LDAP description/info credential harvesting** NEW 65

Plaintext credentials stored in account `description` or `info` LDAP attributes are now a validated, end-to-end attack step: ADscan reads the attribute, recovers the credential, verifies the login, and traces the resulting route across forest boundaries.

— Full mechanism described, no direct command given v11.3.0

06 **Cross-forest attack path graph analysis** NEW 60

Attack paths now cross domain and forest boundaries in multi-domain workspaces, analysed as a unified graph, including cross-forest Kerberoasting, AS-REP Roasting, and ADCS ESC8/ESC11 relay routes across trusts.

— Names specific techniques and route types but no command to invoke v11.3.0

THINNER COVERAGE BELOW

07 **krbtgt-only DCSync mode** NEW 50

DCSync now supports a krbtgt only replication mode that captures the golden-ticket key without pulling every account, offering a minimal-footprint path to proving full domain compromise.

— Mechanism described but no flag or command given v11.3.0

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

Open-source AI penetration testing tool to find and fix your app's vulnerabilities.

> Strix v1.6.0 adds MCP server support with per-task isolation, a managed cloud CLI, workspace file injection into the sandbox, contextual CVSS dependency vulnerability reports, and a wave of new security skills including OWASP LLM Top 10 2026 coverage.

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

01 Workspace file injection via `--workspace-file` NEW

90

Adds a `--workspace-file` CLI flag and `extra_files` config key that place caller-provided files into the sandbox workspace before a scan starts, letting agents reference fixtures or seed data without baking them into the target directory.

Inject a custom config or seed file into the sandbox so agents can reference it during a pentest without baking it into the target directory.

```
$ strix --target ./my-app --workspace-file ./fixtures/test-users.json
```

— Names exact flag, config key, and a runnable command. v1.6.0

02 MCP server support with per-task isolation NEW

75

Adds MCP server support with connections isolated per task, on-demand tool resolution instead of bulk registration, connection status surfaced in the UI, and resilience so connections survive transient transport failures.

— Describes mechanism and behaviour but no concrete command. v1.6.0

03 `--host` option for `strix view` NEW

75

Adds a `--host` option to `strix view` to expose the local results viewer on a custom interface, e.g. `--host 0.0.0.0`.

Expose the scan results viewer on all interfaces at a fixed port when reviewing findings on a remote dev box over SSH.

```
$ strix view my-run-name --host 0.0.0.0 --port 8080 --no-open
```

— Exact flag with a runnable example command. v1.6.0

04 **`update_vulnerability_report` agent API** NEW 60

Adds an `update_vulnerability_report` agent API so an agent can revise an already-filed finding during a run.

— Names the API but no usage detail. v1.6.0

THINNER COVERAGE BELOW

05 **`strix cloud` subcommand for managed platform** NEW 55

Adds a `strix cloud` subcommand for the managed platform CLI, covering login, scans, billing, and the full cloud API.

— Names subcommand and scope, no example usage. v1.6.0

06 **Contextual CVSS breakdown for dependency reports** IMPROVED 55

Adds contextual CVSS breakdown and required usage evidence on dependency vulnerability reports.

— Describes what changed but no interface detail. v1.6.0

07 **New security skill coverage across domains** NEW 50

Adds OWASP LLM Top 10 2026 skill coverage, an Azure and Entra security skill, an argument injection security skill, ecosystem supply-chain security skills, Hurl and Hypothesis security playbooks, and semantic browser and Electron security skills.

— Names six skill areas but no mechanism given. v1.6.0

08 **Faster startup and scan launch** IMPROVED 40

Bootstraps Caido concurrently with scan start to reduce scan startup latency, and moves heavy imports off the startup path, pre-warming them in the background to reduce CLI startup time.

— Describes two startup optimizations without measurable numbers. v1.6.0

09 **CSV injection hardening in report output** IMPROVED 35

Hardens CSV injection in report output.

— Brief security fix note, no mechanism given. v1.6.0

10 **Coverage as first-class artifact with evidence discipline** IMPROVED 30

Promotes coverage as a first-class artifact with evidence discipline in the agent layer.

— Vague description with no named surface. v1.6.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

Nmap

Feedback ⓘ changes since 2026-08-31 CODE ↗

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap NSE gets a 400 MB memory cap and a 5-second execution watchdog that kills non-yielding threads.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-31
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-31
```

Enforces a 400 MB memory limit on the Nmap Scripting Engine (NSE) Lua runtime to bound per-scan resource consumption.

└──▷ BREAKING ON UPGRADE

! NSE scripts that legitimately run without yielding for more than 5 seconds will now be forcibly terminated, potentially breaking long-running scripts that do not yield control.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

cve-lite-cli

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

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

> cve-lite-cli's sole release window feature is an EPSS Priority Signal that fuses CVSS severity with FIRST.org exploitation-likelihood data into four actionable tiers across JSON, terminal, and HTML outputs.

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

01 **EPSS Priority Signal across JSON, terminal, and HTML output**

NEW 91

Adds an EPSS Priority Signal tier system combining CVSS severity with FIRST.org EPSS exploitation-likelihood percentile into four tiers: `fix_now` (high impact + top 10% EPSS), `fix_soon` (high impact), `monitor` (lower severity + top 10% EPSS), and `low_priority` (all else). JSON output gains `epssScores` and `prioritySignal` fields, set to `null` when no CVE alias resolves against the FIRST.org API; the signal appears as a new column in the verbose terminal table (`--verbose`) with an inline `⚡ Fix Now` flag in compact mode, and as a column plus interactive four-tier legend in HTML reports (`--report`).

Triage a project by exploitation likelihood in CI — pipe JSON output and filter for packages the model rates as needing immediate action.

```
$ cve-lite /path/to/project --json | jq '[.[] |
select(.prioritySignal == "fix_now")]'
```

— Names exact fields, tiers, flags, and a runnable CI filter example v1.33.0

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ 6 RELEASES · 2026-09-01 BLOG ↗

Join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> GitHub's biggest ships this window are file attachments landing in the CLI, Copilot code review gaining the power to formally approve pull requests, and GA launches for Enterprise Live Migrations and the Claude Fable 5.1 model.

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

01 File attachments via `--attach` in GitHub CLI NEW

95

Adds a repeatable `--attach` flag to `gh issue create`, `gh issue edit`, `gh issue comment`, `gh pr create`, `gh pr edit`, and `gh pr comment`, uploading a local file and embedding it inline in the Markdown body. Supports PNG, JPEG, GIF, WebP, SVG, MP4, MOV, and WebM formats, with size limits of 10 MB for images/GIFs, 10 MB for video on Free plans, and 100 MB for video on paid plans. Inline alt text can be set via `#` syntax, e.g. `--attach './login.png#The login error state'` (falls back to filename if omitted); local image paths already referenced in the Markdown body are rewritten to the uploaded asset URL with alt text preserved, and unreferenced attachments are appended at the end.

File a bug report with a screenshot attached in one command so the issue shows the actual problem immediately.

```
$ gh issue create --title 'Login button misaligned on mobile' --body 'See screenshot for the rendering issue.' --attach './screenshot.png#Login button misaligned'
```

Open a PR with before/after video clips so reviewers see the visual change without pulling the branch.

```
$ gh pr create --title 'Smooth out modal animation' --body 'Before and after recordings attached.' --attach './before.mp4#Before' --attach './after.mp4#After'
```

Add a clarifying screenshot to an existing issue comment when a text description alone is ambiguous.

```
$ gh issue comment 42 --body 'Reproduced on v2.99.0 — error visible here:' --attach './error.png#Error on screen'
```

— Exact flag, formats, size limits and runnable commands given.

GitHub CLI: Media in issues, pull reque…

02 **Enterprise Live Migrations GA via `gh elm`** NEW 90

Enterprise Live Migrations (ELM) is now generally available as a GHES-appliance service for near-zero-downtime repository migration to GitHub Enterprise Cloud with Data Residency (GHEC DR), supporting GHES versions 3.17.18+, 3.18.12+, 3.19.9+, 3.20.3+, 3.21.3+, and 3.22.0+. The `gh elm` GitHub CLI extension manages the full lifecycle — creating migrations, monitoring sync progress (with `--json` output for scripting), and initiating cutover — against the GHES REST API, with resource-level progress tracking that surfaces failures before cutover. ELM can run concurrently with GitHub Enterprise Importer (GEI), letting teams pick the right tool per repository.

— Names CLI extension, API and supported versions but no sample command.

Enterprise Live Migrations from GHES to…

03 **Expiration dates for individual user budgets** NEW 90

The Budgets REST API adds an `expires_at` field when creating or updating an individual user budget, supporting no expiration (default), start of next billing cycle, or a specific date; the same options are exposed in the billing settings UI and can be changed or cleared at any time. Expired budgets are automatically removed on their expiration date, falling the user back to their cost center per-user budget or the universal budget. Available on Copilot Business and Copilot Enterprise plans.

Create a new individual user budget with an expiration at the start of the next billing cycle directly from billing settings.

📍 1. Go to your enterprise Billing settings › Budgets. 2. Click `Create budget` or select an existing individual user budget to edit. 3. Under `Expiration`, choose 'Expiration at the start of the next billing cycle'. 4. Save the budget — GitHub will automatically remove it on that date.

— Names the API field, UI path and exact expiration options.

Set an expiration date for individual u…

04 **Claude Fable 5.1 model in Copilot** BREAKING 90

Claude Fable 5.1 joins the Copilot model picker across Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, github.com, GitHub Mobile, JetBrains IDEs, Xcode, and Eclipse, available to Copilot Pro+, Max, Business, and Enterprise plans. Enterprise/Business admins enable it via the `Claude Fable 5.1` policy in Copilot settings (off by default), and eligible enterprise customers can get zero data retention for Claude Fable 5.1 and Claude Fable 5 through a time-bound exemption while Anthropic rolls out Enterprise Frontier Safeguards. Unlike every other Claude model in Copilot, which runs under zero data retention (ZDR) by default, Claude Fable 5.1 and Claude Fable 5 retain prompts and outputs by default to operate safety classifiers, and enabling the policy constitutes acknowledgment of this requirement.



Enable Claude Fable 5.1 for your organization as a Copilot Enterprise or Business admin so team members can select it in the model picker.

> 📍 In GitHub Copilot settings for your organization, navigate to Policies, locate the `Claude Fable 5.1` policy, and toggle it on. Note: enabling constitutes acknowledgment of Anthropic's data retention requirement for Mythos-class models.

— Names the policy, plans, and default-retention behavior change.

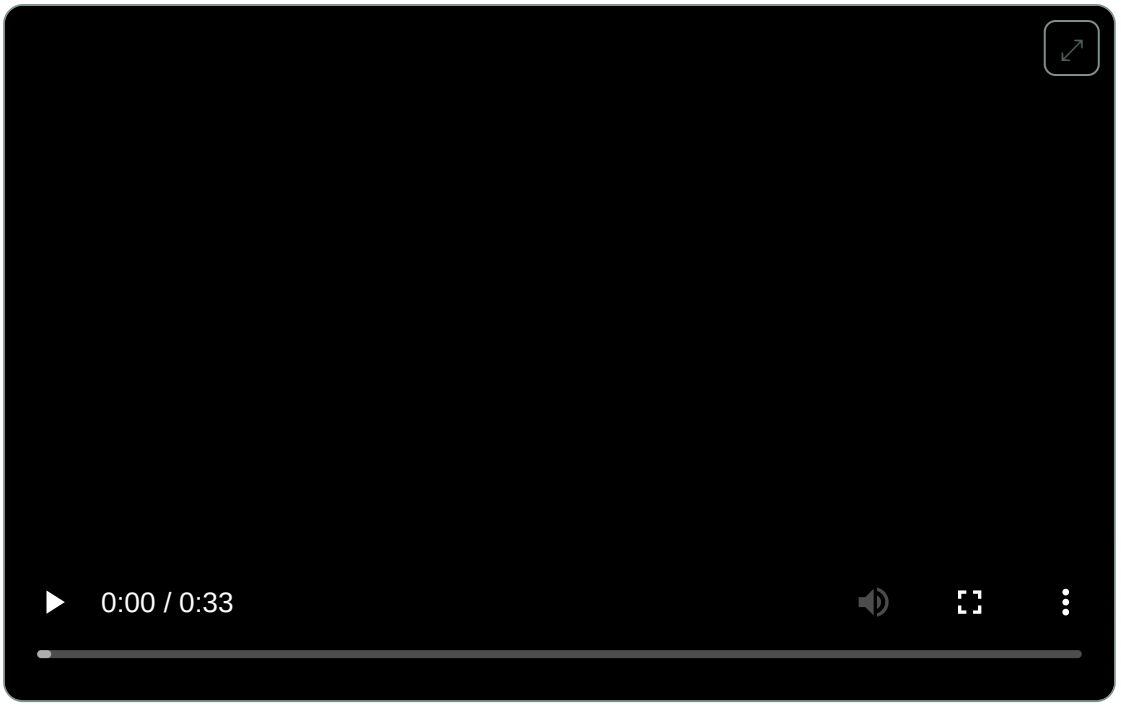
`Claude Fable 5.1 is generally available…`

05 Copilot code review can approve pull requests NEW

85

Every Copilot code review now includes an approval assessment in its overview comment signaling whether Copilot considers the PR ready to approve (this alone doesn't count toward merge requirements). When enabled by an admin (off by default, configurable at enterprise, organization, and repository level, including per-file-path controls at the repo level), Copilot can submit a formal approval that counts toward a

repository's required-approvals rule; the approval is automatically dismissed when new commits are pushed, matching human-reviewer behavior.



Enable Copilot approvals for specific repositories in your organization so Copilot-issued approvals count toward required-reviewer rules.

> 📍 In your organization settings, go to Copilot › Code review › Approvals, set the policy to `Enable for specific repositories`, and select the target repositories. At the repository level, go to Settings › Copilot › Code review › Approvals, turn approvals on, and optionally restrict which file paths Copilot is allowed to approve.

— Describes mechanism and scope with a clear settings path.

`Copilot code review can now approve pul…`

THINNER COVERAGE BELOW

06 **Block users from discussion comments** NEW 55

Adds in-context block/unblock controls to discussion comments in personal repositories, matching controls already available on issue and pull request comments, accessible via the `More` menu on any comment with an optional private note.

Block a disruptive commenter without leaving the discussion thread in your personal repository.

> 📍 In a discussion comment on your personal repository, open the More menu (...) on the comment › select `Block user` › optionally enter a private note explaining the reason › confirm your selection.

— Simple UI addition, clear navigation but little mechanism.

`Block users from discussion comments in…`

└──▷ **BREAKING ON UPGRADE**

! Claude Fable 5.1 (and Claude Fable 5) require data retention by default — prompts and outputs are retained by Anthropic to operate safety classifiers — unlike all other Claude models in Copilot, which continue to operate under zero data retention (ZDR).

WAS THIS USEFUL?  

aflplusplus

Feedback ⓘ 1 RELEASE · 2026-09-02 NOTES ↗

The fuzzer afl++ is afl with community patches, qemu 5.1 upgrade, collision-free coverage, enhanced laf-intel & redqueen, AFLfast++ power schedules, MOpt mutators, unicorn_mode, and a lot more!

> AFL++ v5.03c's headline additions are Value Profile-guided mutation, four new LLVM instrumentation tuning env vars, and a starved-mode fuzzing state with automatic queue minimization, alongside a breaking switch to POSIX shared memory on macOS and removal of classic AFL instrumentation and the afl-as wrapper.

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

01 New LLVM instrumentation coverage tuning options NEW

87

Adds four env vars to tune LLVM instrumentation density and sensitivity: `AFL_LLVM_DENSE=1` disables PCGUARD basic block pruning for denser coverage maps; `AFL_LLVM_MINMAX=1` scores min/max/abs intrinsics (branchless clamps), adding ~10% map entries at a ~20% throughput cost, default off; `AFL_LLVM_FUSED=1` scores both halves of fused conditions (e.g. if (a && b) compiled to a single `and i1`); `AFL_LLVM_VECTORS=1` instruments vector selects and vector min/max with one guard pair per lane, default off.

Score branchless min/max/abs clamps and fused conditions for deeper coverage of optimizer-transformed code, accepting the throughput trade-off.

```
$ AFL_LLVM_MINMAX=1 AFL_LLVM_FUSED=1 afl-clang-fast -o target
target.c
afl-fuzz -i seeds/ -o out/ -- ./target @@
```

— Names all four env vars with exact behaviour and cost trade-offs. v5.03c

02 Value Profile fuzzing mode NEW

86

Adds `AFL_LLVM_VALUE_PROFILE=1` to enable Value Profile instrumentation in fuzz targets, paired with a new `afl-fuzz` flag `-r <seconds>` that activates value-profile-guided mutation after a delay, letting a baseline fuzzing period run before switching to value-profile-informed mutation.

Enable Value Profile instrumentation to capture comparison operand diversity, then activate it in afl-fuzz after 30 seconds of baseline fuzzing.

```
$ AFL_LLVM_VALUE_PROFILE=1 afl-clang-fast -o target target.c
afl-fuzz -i seeds/ -o out/ -r 30 -- ./target @@
```

— Names exact env var and paired flag with a runnable example. v5.03c

03 **POSIX shared memory on macOS** BREAKING 81

macOS builds now use POSIX shared memory instead of SysV, removing the `kern.sysv.shmseg` / `kern.sysv.shmmax` ceiling constraints — CmpLog alone requires a 145 MB segment, far above the 4 MB macOS default. Targets built with an older `afl-cc` still expect a SysV id in `__AFL_SHM_ID` and must be recompiled with the new compiler to work.

— Explains mechanism, concrete limits, and migration requirement. v5.03c

04 **Starved mode with automatic queue minimization** NEW 79

afl-fuzz now implements a `starved` mode: when no new finds occur for an extended period, seed and mutation variability increases automatically (including the splicing phase), visible in the UI. The new `AFL_STARVED_MINIMIZE_QUEUE` env var makes starved mode minimize the queue (like `afl-cmin`) when it fails to find new coverage, and removed entries' coverage becomes rediscoverable.

When a long-running campaign stalls, trigger queue minimization automatically so stale entries are pruned and their coverage becomes rediscoverable.

```
$ AFL_STARVED_MINIMIZE_QUEUE=1 afl-fuzz -i seeds/ -o out/ --
./target @@
```

— Describes trigger, mechanism and named env var with example. v5.03c

05 **Symlink and foreign sync directory handling** IMPROVED 65

`afl-fuzz` `-i` input directory and `afl-showmap` `-i` now follow symlinked test cases and symlinked subdirectories; foreign sync directories (`-F`) also import symlinked test cases. Unreadable foreign sync (`-F`) directories are now reported once per directory and listed at startup, replacing silent failure.

— Names affected flags and behaviour but no example command. v5.03c

06 **afl-cmin default implementation switch** IMPROVED 62

Switches `afl-cmin.c` to be the default `afl-cmin` implementation; `-T all` on both the C and Python variants now counts CPUs the process is actually allowed to run on, rather than all system CPUs.

— Explains the CPU-counting behaviour change but gives no example. v5.03c

07 **Trimming stage saves crashes and hangs** IMPROVED 62

The `afl-fuzz` trimming stage now saves crashes, hangs, and new coverage discovered while shortening inputs, visible as `op:trim` / `op:ptrim` in output file names.

— Names exact output markers for the behaviour change. v5.03c

THINNER COVERAGE BELOW

08 **Removal of legacy AFL instrumentation and afl-as** BREAKING 59

Removes classic AFL instrumentation (colliding coverage) from `afl-cc`, superseded by `AFL_LLVM_PATH` and `AFL_LLVM_CALLER`; also removes the obsolete `afl-as` assembler wrapper.

— Names removed components and their replacement, no migration steps. v5.03c

09 **SIGUSR2 forces immediate sync** NEW 55

afl-fuzz now handles `SIGUSR2` to force an immediate sync (both AFL's internal sync and foreign `-F` sync directories) once the current queue entry finishes processing.

— Names signal and behaviour but gives no usage example. v5.03c

10 **afl-merge corpus merging tool** NEW 45

Adds `afl-merge` as a symlink tool to merge new files into an existing (non-minimized) corpus.

— Names tool and purpose but no usage example. v5.03c

11 **AFL_DUMP_MAP_SIZE diagnostic env var** NEW 37

Adds `AFL_DUMP_MAP_SIZE` to print the coverage/IJON breakdown to stderr.

— Bare description of a single env var, no example. v5.03c

12 **Exec-count-based mode switching** IMPROVED 30

Mode-switching limits in `afl-fuzz` are now based on execution counts rather than elapsed time.

— States the change with no further mechanism detail. v5.03c

└──▷ BREAKING ON UPGRADE

! macOS targets must be recompiled with the new `afl-cc`: a target built by an older `afl-cc` still expects a SysV id in `__AFL_SHM_ID` and will not work with the new POSIX shared memory.

! Classic AFL instrumentation (colliding coverage) has been removed from `afl-cc`; projects relying on it must switch to `AFL_LLVM_PATH` or `AFL_LLVM_CALLER`.

! The obsolete `afl-as` assembler wrapper has been removed.

WAS THIS USEFUL? 👍 👎

JFrog

Feedback

3 RELEASES · 2026-09-02

BLOG ↗

JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

> JFrog expanded AppTrust with auto-assembled, signed release-evidence chains and out-of-the-box NIST SSDF / EU CRA compliance frameworks, while JFrog AI Catalog grew into a fuller AI control plane with governed registries for MCP servers, skills and plugins, Agent Packages, semantic scanning, Shadow AI Detection, and runtime policy enforcement via Agent Guard.

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

01 **Agent Guard runtime policy enforcement** NEW 75

Agent Guard, installed via the official JFrog agent plugin, enforces organizational policy at runtime by checking every MCP tool call and skill call against defined policies, regardless of agent harness or coding environment. Setup: install the JFrog agent plugin under AI Catalog > Agent Plugins, configure the MCP Registry's allowed servers and restricted tool calls per project, then enable Agent Guard in plugin settings.

Enforce that a coding agent in Cursor or Claude Code can only install MCP servers, skills, and plugins that have passed your organization's security checks via Agent Guard at runtime.

> 📍 Steps: 1. In the JFrog platform, go to AI Catalog › Agent Plugins and install the JFrog agent plugin for your coding environment (Cursor, Claude Code, VS Code, OpenCode, or Codex). 2. Configure your MCP Registry: define the allowed MCP servers per project and specify any restricted tool calls. 3. Enable Agent Guard within the plugin settings to activate runtime policy enforcement for MCP tool calls and skill calls.

— Concrete setup steps and enforcement scope given.

The Evolution of JFrog AI Catalog: Your…

02 **AppTrust Prompt-to-Release Traceability** NEW 70

AppTrust now auto-assembles a signed audit evidence chain — commits, PRs, Jira ticket transitions, and approver records — and attaches it to the application version before any promotion gate runs. It validates full Jira ticket transition history (creation → in-progress → done, with dates) rather than just checking ticket existence, and detects self-approval violations by mapping each PR's merger and approver back to the committer. Incoming JFrog Agent Plugins will capture AI agent session evidence (prompts, code produced, and a full session bill of materials including which MCPs, skills, and agentic assets were used) as signed evidence stored in Artifactory. Git-based evidence is generally available

now; agent session evidence and the session bill of materials are planned for later this year.

— Rich mechanism and rollout timeline, but no exact command or config key.

Every Release Needs a Chain of Custody.…

03 **Out-of-the-Box Compliance Frameworks in AppTrust** NEW

70

Launches with NIST SSDF and the EU Cyber Resilience Act pre-mapped to Policy-as-Code rules, requiring no Rego to write and no manual control mapping. A compliance catalog inside AppTrust lets AppSec teams pick a framework and see every control translated into plain requirements linked to specific Policy-as-Code rules enforced at release gates, with uncovered controls flagged and satisfied controls shown as covered by default. Full SDLC traceability (agent intent sessions, Git commits, PRs, Jira tickets) feeds live data to these rules. Framework configuration is split from rollout control: AppSec owns which controls/rules are active, while development managers control application scope, whether rules run as warnings or hard failures, and the warning-to-block conversion window. Available now as part of the Ultimate Security Bundle.

— Names frameworks and bundle, but no runnable command.

Every New Compliance Framework Restarts… Every Release Needs a Chain of Custody.…

04 **MCP, Agent Skills, and Agent Plugins Registries in AI Catalog** NEW

70

Adds an MCP Registry to define which MCP servers are allowed per project, restrict specific tool calls, and scan local MCP packages before an agent can use them; an Agent Skills Registry to curate, sign, verify, and scan packaged agent behaviors before execution; and an Agent Plugins Registry to control, sign, verify, and scan plugin bundles agents install. Models, MCP servers, skills, plugins, and agent packages are all stored as versioned, traceable artifacts in Artifactory under a single RBAC model and audit trail.

— Each registry's mechanism named; no exact command given.

The Evolution of JFrog AI Catalog: Your…

05 **Shadow AI Detection in AI Catalog** NEW

65

Automatically detects and categorizes every AI asset across an environment as managed, unmanaged, partially managed, or malicious, via AI Catalog > Shadow AI Detection; unvetted assets can then be promoted into the appropriate governed registry (MCP Registry, Agent Skills Registry, or Agent Plugins Registry).

Detect unvetted AI assets sprawling across your environment before they reach an agent, by reviewing Shadow AI Detection results categorized as managed, unmanaged, partially managed, or malicious.

📍 Steps: 1. In the JFrog platform, navigate to AI Catalog › Shadow AI Detection. 2. Review the categorized inventory of AI assets across your environment (managed, unmanaged, partially managed, malicious). 3. Promote assets you wish to approve into the appropriate governed registry (MCP Registry, Agent Skills Registry, or Agent Plugins Registry).

— Exact UI navigation and categories given as a usable path.

The Evolution of JFrog AI Catalog: Your…

06 **AI Coding Agent Plugins for Claude Code, Cursor, VS Code, OpenCode, and Codex**

NEW

60

The AI Coding Agent Plugin now supports Claude Code, Cursor, VS Code, and OpenCode as agents, and JFrog released official agent plugins for Claude Code, Cursor, VS Code, OpenCode, and Codex, installed via AI Catalog > Agent Plugins, to bring organizational policy control directly into the coding agent environment.

— Names agents and install path, no further mechanism detail.

The Evolution of JFrog AI Catalog: Your…

THINNER COVERAGE BELOW

07 **Agent Packages support (APM standard)** NEW

45

Adds Agent Packages support to ship an agent's entire context and dependency tree as a single versioned, resolvable package compliant with the open-source APM (Agent Package Manager) standard.

— Names the APM standard but no usage steps. The Evolution of JFrog AI Catalog: Your…

08 **Semantic scanning of AI assets** NEW

40

Introduces semantic scans that analyze what an asset instructs an agent to do, rather than matching against known-bad hashes, in order to catch novel malicious skills and MCP servers.

— Explains the mechanism but no configuration surface.

The Evolution of JFrog AI Catalog: Your…

WAS THIS USEFUL? 👍 👎

Snyk

Feedback

1 RELEASE · 2026-08-31

BLOG ↗

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

> Snyk shipped Repo Content Sync, now generally available for Enterprise plans, which automatically keeps Projects in sync with manifest, Dockerfile, and IaC files as they're added, moved, or removed from monitored branches.

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

01 **Repo Content Sync for automatic Project management** NEW

74

Repo Content Sync (GA for Enterprise plans) automatically creates Projects for newly merged manifest, Dockerfile, and IaC configuration files, and deactivates Projects whose files are removed or renamed, eliminating manual re-import. It is on by default for Enterprise organizations with no configuration required; organizations already opted in via Snyk Preview retain their existing setting. To permanently suppress a deleted Project from being recreated when its source file still exists, open the Project, click Delete, and select 'exclude it from future scans' in the delete dialog before confirming.

Permanently suppress a deleted Project from being recreated by sync when its source file still exists in the repo.

> 📍 In the Snyk UI, open the Project you want to remove, click Delete, and in the delete dialog select the option 'exclude it from future scans' before confirming.

— Names mechanism, default behavior, and exact UI steps to opt out. 2026-08-31-baf195e9

WAS THIS USEFUL? 👍 👎

Snyk Code

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Snyk Code is a developer security tool that scans source code to identify and fix vulnerabilities, open source issues, and code quality problems.

> Snyk Code's repo content sync now covers all Snyk project types and works beyond default branches.

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

01 **Expanded repo content sync coverage** IMPROVED 50

Repo content sync now covers Code, open source, secrets, infrastructure as code, and container (Dockerfile) Projects — not just a subset of project types — and works on custom branches in addition to default branches.

— Names project types and branch scope but no config or command shown.

2026-08-31-b80221ae

WAS THIS USEFUL? 👍 👎

Snyk Open Source

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Snyk Open Source scans open-source dependencies for known vulnerabilities and license issues, helping developers find and fix risks in their software supply chain.

> Snyk Open Source's repo content sync now extends coverage across all project types and to custom branches, not just default branches.

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

01 **Expanded repo content sync coverage** IMPROVED 48

Repo content sync now covers Code, open source, secrets, infrastructure as code, and container (Dockerfile) project types — not just a subset — and now works on custom branches in addition to default branches.

— Names project types and branch scope but no config/API detail. 2026-08-31-b80221ae

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

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

> Infisical's v0.165.1 release focuses on PKI expansion (GCP Certificate Manager sync, agent renewal tracking, pre-fetch commands), tighter access control (folder RBAC, ABAC `trim suffix`, Kubernetes MI auth), and several smaller usability and provider tweaks across secret management and PAM.

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

01 **GCP Certificate Manager Sync for PKI** NEW 60

Certificates managed in Infisical's PKI can now be synced to Google Cloud Certificate Manager, letting teams keep GCP-hosted certificate resources in step with Infisical as the source of truth.

— Names target integration but no config steps given v0.165.1

02 **`trim suffix` function in ABAC policies** NEW 60

Adds a `trim suffix` function usable within attribute-based access control (ABAC) policy expressions.

— Named function directly usable in policy expressions v0.165.1

THINNER COVERAGE BELOW

03 **Folder-level RBAC** NEW 50

Role-based access control can now be scoped at the folder level within projects, rather than only at the project level.

— Clear scope but no policy syntax shown v0.165.1

04 **Webhooks for approval and change request events** NEW 50

Infisical now emits webhook events for approval and change request workflows, allowing external systems to react to secret approval lifecycle events.

— Describes event type but no payload or endpoint given v0.165.1

05 **Kubernetes Managed Identity auth template** NEW 50

Adds a Kubernetes Managed Identity (MI) auth template for machine identity authentication.

– Names auth method but no setup detail v0.165.1

06 **Port configuration for Splunk dynamic secrets** IMPROVED 50

Adds port number configuration to the Splunk dynamic secrets provider.

– Names provider and setting but no field name shown v0.165.1

07 **Certificate renewal visibility in PKI agent** NEW 45

The PKI agent now exposes the latest renewal of a certificate so agents can detect and follow certificate renewals.

– Explains behavior but no API/field name given v0.165.1

08 **Default interactive shell for PAM SSH access** IMPROVED 40

SSH access in Privileged Access Management (PAM) now defaults to an interactive shell.

– Clear default change, no config override mentioned v0.165.1

09 **Pre-fetch commands for PKI sync agents** NEW 35

Adds pre-fetch commands for PKI sync agents on Windows and Linux platforms.

– Bare mention, no command syntax or use case v0.165.1

10 **Project-level secret retention limits** NEW 35

Adds the ability to set secret retention limits at the project level.

– No limit values or config path specified v0.165.1

11 **Revamped replicate secrets workflow** IMPROVED 35

The replicate secrets workflow in the secret manager UI has been revamped.

– No specifics on what changed in the workflow v0.165.1

12 **Select-all for Vault import secret path picker** IMPROVED 30

Adds a select all option for the Vault import secret path Combobox in the UI.

– Small UI addition, exact location given v0.165.1

13 **DigiCert profile helper text** IMPROVED 25

Adds helper text for DigiCert profiles during certificate issuance.

— Minor UI copy change with no further detail v0.165.1

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Mobile Verification Toolkit

Feedback ⓘ

1 RELEASE · 2026-09-01 NOTES ↗ CODE ↗

MVT (Mobile Verification Toolkit) helps with conducting forensics of mobile devices in order to find signs of a potential compromise.

> MVT adds a plugin system for extending forensic analysis, a unified `mvt` CLI with shell completion, and substantially deeper Android bugreport and ADB forensics.

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

01 Shell completion via `mvt completion` NEW

90

Adds `mvt completion` to generate shell completion scripts for Bash, Zsh, and Fish across `mvt`, `mvt-ios`, and `mvt-android`; pass the shell name (e.g. `mvt completion bash`) to emit the script directly, and `--install` to write it.

Generate and install shell completion for Bash so all three MVT commands are tab-completable in your daily workflow.

```
$ mvt completion bash --install
```

▶ View command output v2026.9.1

— Exact runnable command with flags given v2026.9.1

02 Plugin system for custom modules and commands NEW

88

Adds a plugin system: third-party packages can register forensic modules and top-level CLI commands via Python entry points, and custom modules can replace built-in MVT modules. `mvt plugins` lists installed plugins and checks for updates, and `mvt.plugin` is now a public Python import surface for writing plugin packages.

Check what plugins are installed and whether any have updates available, before running an analysis.

```
$ mvt plugins
```

— Names the command and import surface with runnable example v2026.9.1

03 **Expanded Android bugreport parsing** IMPROVED 80

Preserves Android tombstone abort messages and protobuf `causes` lists in both text-parsed and protobuf-parsed output; parses package installer attribution, process tables, system settings, and mountinfo from bugreports; and adds parsing of system package flags, accessibility service states per user, ADB binary XML keys, AppOps UID and event details, battery history events, database operation details, multi-user package details, platform compatibility overrides, and multiline Android properties.

— Enumerates every new bugreport data type parsed v2026.9.1

04 **ADB trusted key alerting in AndroidQF** NEW 78

Alerts on unexpected trusted ADB keys in AndroidQF acquisitions: raises LOW-severity alerts for malformed keys, keys from a different host, and trusted keys when no acquisition host key is present, plus an informational alert when the acquisition host key was last connected more than one day before collection.

— Detailed alert conditions but no command to trigger check v2026.9.1

05 **Top-level `mvt` command** NEW 70

Adds a platform-neutral `mvt` command hosting `version`, `completion`, `plugins`, and `download-iocs` subcommands; running `mvt` alone shows the installed version, update notices, and available commands.

— Names subcommands but no full usage walkthrough v2026.9.1

THINNER COVERAGE BELOW

06 **Encrypted backup.ab handling in AndroidQF** IMPROVED 55

AndroidQF runs no longer abort when encountering an encrypted `backup.ab`; the encrypted backup is skipped and analysis continues.

— Names the file and behavior change, no further steps v2026.9.1

07 **Modules skip on missing dependencies** IMPROVED 35

Modules with unavailable dependencies are now skipped instead of aborting the entire run.

— Behavior described but no mechanism or affected module list v2026.9.1

08 **WhatsApp contacts forensic module** NEW 35

Adds a WhatsApp contacts forensic module.

— Bare mention, no detail on what it extracts v2026.9.1

09 **ARM64 Docker images** NEW 35

Builds and publishes Docker images for ARM64.

— Bare announcement with no pull instructions v2026.9.1

10 **Timezone provenance on file timestamps** IMPROVED 25

Adds timezone provenance to file timestamps in analysis output.

— Very brief, no example of output change v2026.9.1

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗

Fastly is a content delivery network and edge computing platform that accelerates website and application performance globally.

> Fastly shipped two new platform capabilities—schema-based API request validation and on-platform log storage/monitoring—plus a new client-disconnect metric across its analytics APIs.

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

01 **status_499 metric for client disconnects** NEW 75

Adds a `status_499` metric to the Historical Stats API, Real-Time Analytics API, Domain Inspector Historical API, and Domain Inspector Real-Time API to surface client disconnects from Varnish services.

– Names four specific APIs and the exact metric field

Add status 499 metric to Historical Sta…

THINNER COVERAGE BELOW

02 **API Enforcement for schema-based request validation** NEW 40

A new API Enforcement feature validates incoming API requests against user-defined schemas, and can be configured to either log or block requests that don't conform to the schema.

– Describes behavior but no config keys or endpoints named product docs

03 **Log Explorer & Insights for on-platform logging** NEW 40

A new Log Explorer & Insights feature enables proactive storage, inspection, and monitoring of log data directly on the Fastly Platform, without needing to ship logs to a third-party system.

– Explains purpose but lacks specific UI paths or API details product docs

WAS THIS USEFUL? 👍 👎

Sophos Email Security

Feedback ⓘ 1 RELEASE · 2026-08-31

NOTES ↗

Sophos Email Security is a cloud-based email protection solution that detects and blocks malware, phishing, spam, and advanced threats.

> Sophos Email Security added recipient notifications for Data Control rule actions and automatic Intelix-based release of cleared quarantined emails.

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

01 **Automatic Intelix-based quarantine release** NEW 65

Automatically releases quarantined emails when an Intelix scan determines the email is not malicious or suspicious, reducing manual triage of held mail. Enabled via Scan attachments with Intelix automatic release option in an Email Security > Policies > Data Control rule set to Quarantine.

Configure automatic quarantine release so emails cleared by Intelix are freed without manual intervention — reducing analyst queue depth for low-risk holds.

📍 In the Sophos Central console, go to Email Security › Policies › Data Control, open or create a rule, set the action to Quarantine, then under Scan attachments with Intelix enable automatic release so emails assessed as not malicious or suspicious are released automatically. Under Notifications, add the recipient addresses that should be alerted when the rule triggers.

— Names exact console path and toggle for setup 2026.35Released August 31, 2026

THINNER COVERAGE BELOW

02 **Notification recipients for Data Control rules** NEW 50

Enables configuration of notification recipients for Data Control rule actions, so the right people are alerted when a Data Control policy triggers. Configured under Email Security > Policies > Data Control, in a rule's Notifications section.

Configure automatic quarantine release so emails cleared by Intelix are freed without manual intervention — reducing analyst queue depth for low-risk holds.

📍 In the Sophos Central console, go to Email Security › Policies › Data Control, open or create a rule, set the action to Quarantine, then under Scan attachments with Intelix enable automatic release so emails assessed as not malicious or suspicious are released automatically. Under Notifications, add the recipient addresses that should be alerted when the rule triggers.

— Names UI path but no deeper mechanism 2026.35Released August 31, 2026

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ 1 RELEASE · 2026-09-02 CHANGELOG ↗

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

> Chainguard's latest changelog introduces open-beta custom container pull policies and enables cross-language build pinning for Chainguard Libraries, alongside a batch of new container images and extended support windows.

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

01 **Extended support for Istio 1.29 and Teleport 18 images**

IMPROVED 58

Extends the active support window for Istio 1.29 images (`istio` , `istio-cni` , `istio-envoy` , `istio-operator` , `istio-pilot-agent` , `istio-pilot-discovery` , `ztunnel`) and `teleport:18` through 2027-02-28.

— Gives exact images and end date but no other mechanism changelog-20260902-4bc59dde

WAS THIS USEFUL? 👍 👎

kubescape

Feedback

1 RELEASE · 2026-09-02

NOTES ↗

Kubescape is an open-source Kubernetes security platform for your IDE, CI/CD pipelines, and clusters. It includes risk analysis, security, compliance, and misconfiguration scanning, saving Kubernetes users and administrators precious time, effort, and resources.

> Kubescape v4.0.13 adds custom Rego rule loading and fully offline image scanning, expands `kubescape fix` to handle cluster reports and JSON manifests with per-control selection, introduces MCP server tools for AI assistants, and ships new RBAC privilege-escalation controls alongside VEX, GitLab and Harbor scan-output support.

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

01 kubescape fix control selection and manifest support IMPROVED

93

`kubescape fix` gains `--include-controls` and `--skip-controls` to choose which controls get remediated, now supports printing patched manifests from cluster scan reports (previously only file-based manifests were supported), and can apply fixes to JSON manifests in addition to YAML.

Auto-remediate only specific controls to limit the blast radius of a fix run.

```
$ kubescape fix --include-controls C-0016,C-0044
/path/to/manifests/
```

— Named flags, before/after, and a runnable example v4.0.13

02 Fully offline image vulnerability scanning NEW

80

Adds `--skip-db-update` to `kubescape scan image` so image vulnerability scans can run in fully offline or air-gapped environments without updating the vulnerability database.

Scan a container image for vulnerabilities in an air-gapped environment without updating the vulnerability DB.

```
$ kubescape scan image myregistry.local/myapp:v1.2.3 --skip-db-
update
```

— Exact flag, subcommand and a runnable example v4.0.13

03 New RBAC privilege-escalation controls NEW

78

New control `C-0298` detects RBAC subjects that can attach ephemeral containers to pods. Adds RBAC-police rules covering `impersonate`, `token-mint`, `escalate`, `bind`, `node-proxy`, and PV create privilege escalation paths.

— Named control ID and six specific privilege paths v4.0.13

04 **Custom Rego rules for scans** NEW 70

Adds a `--custom-rules` flag to load user-authored Rego rules from the standard rule directory layout during scans, letting users extend Kubescape's built-in control set.

— Names the flag but not the directory layout format v4.0.13

05 **Offline evaluation of VAP matchConditions** IMPROVED 65

Evaluates VAP `matchConditions` offline instead of refusing the control, and honors `spec.failurePolicy` for validation expression errors.

— Names fields and behavior change, no user-facing flag v4.0.13

06 **MCP server tools for AI assistants** NEW 60

Adds `scan_controls`, `scan_local_iac_controls`, `list_frameworks`, and `list_controls` tools to the MCP server, exposing Kubescape scanning and control listing to AI assistant integrations.

— Names all four tools but not invocation details v4.0.13

07 **External exposure modeling via Ingress and Gateway API** IMPROVED 60

The exposure engine now models external exposure via Ingress, Gateway API, and Service type, plus `spec.externalIPs`.

— Names all four exposure surfaces, but detection is automatic v4.0.13

THINNER COVERAGE BELOW

08 **MutatingAdmissionPolicy discovery in admission-policy reconciler** NEW 50

Adds `MutatingAdmissionPolicy` discovery and impact matching to the admission-policy map reconciler.

— Names the resource and reconciler, no user action shown v4.0.13

09 **kubescape list controls filtering and config display** IMPROVED 50

Adds filter support to `kubescape list controls` for filtering by framework and search term, and displays the control configuration a scan evaluates against in `kubescape list`.

— Names subcommand, no exact filter flag given v4.0.13

10 **Framework coverage transparency reporting** IMPROVED 50

Flags frameworks scored 100% purely from irrelevant controls and reports resource kinds that no control examined, improving coverage transparency.

— Explains mechanism, no named flag or report field v4.0.13

11 **Format flag for config view** NEW 45

Adds a `--format` / `-f` flag to `kubescape config view` to control the output format.

— Flag named but no example or format list given v4.0.13

12 **Vulnerability adaptors for GitLab and Harbor** NEW 45

Adds native adaptors that convert image scanning output for GitLab and Harbor vulnerability formats.

— Two named integrations, no output format or flag given v4.0.13

13 **Aggregate summary outputs in kubescape diff** IMPROVED 45

Adds aggregate summary outputs to `kubescape diff`.

— Names subcommand but not the summary format v4.0.13

14 **VEX status in image scan reports** NEW 40

Surfaces VEX (Vulnerability Exploitability eXchange) status directly in image scan reports.

— Names VEX but no mechanism or config detail v4.0.13

15 **Audit logging for posture exceptions** NEW 35

Adds opt-in audit logging for posture exceptions.

— Opt-in noted but no config key or flag named v4.0.13

16 **Multi-arch image scanning support** IMPROVED 30

Scans the correct platform for multi-arch container images instead of a single default architecture.

— Generic description, no configuration surface v4.0.13

17 **URL context in repository scanner errors** IMPROVED 15

Adds URL context to repository scanner error messages for easier diagnosis.

— Minor diagnostic tweak with no actionable detail v4.0.13

WAS THIS USEFUL? 👍 👎

Snyk Container

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗

Snyk Container scans container images and Kubernetes workloads for vulnerabilities and misconfigurations, helping developers fix and secure them.

> Snyk Container now benefits from automatic repository content sync, which keeps Dockerfile-based Projects in step with changes to a repo without manual re-import.

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

01 **Automatic Project sync for Dockerfile changes** NEW 67

Snyk automatically creates and begins monitoring new Projects when Dockerfile or configuration files are added to a repository, deactivates Projects whose source files are removed, and tracks moved or renamed files by activating a Project at the new path while deactivating the one at the old path. Sync support covers Code, open source, secrets, infrastructure as code, and container (Dockerfile) Project types, and works on custom branches in addition to default branches.

— Explains mechanism and scope but no config keys or commands given. 2026-08-31-338a15d8

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 1 RELEASE · 2026-09-02 NOTES ↗ CODE ↗

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

> Prowler shipped 25 new checks spanning AWS and seven other providers, two new compliance frameworks (CIS Google Workspace v1.4.0 and NCSC Cyber Essentials v3.3), Slack as an alert destination, and a browser-based Import Findings workflow, alongside SSRF-safe private-registry scanning for its Image provider and several Jira/MCP automation improvements.

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

01 **Private network allowlist for on-prem registry scanning** NEW

80

Adds `PROWLER_IMAGE_PROVIDER_ALLOWED_PRIVATE_NETWORKS` environment variable accepting a comma-separated list of IPs and CIDRs (e.g. `192.168.65.254/32,10.20.0.0/16`) so the Image provider can reach on-premises registries while keeping link-local, loopback, and other unlisted private addresses blocked by the SSRF guard. Malformed entries cause a startup failure, and any non-empty allowlist is logged as a relaxed security control for audit purposes.

— Names exact env var and CIDR example but no runnable command. 5.41.0

02 **Configurable log group prefixes for AgentCore data protection check**

NEW

80

Adds `agentcore_log_group_name_prefixes` key to `config.yaml` to configure which CloudWatch log group prefixes the `cloudwatch_log_group_agentcore_data_protection_policy_enabled` check evaluates.

Scope the AgentCore log-group data-protection check to your organization's specific log group prefixes.

```yaml
agentcore_log_group_name_prefixes:
  - "/aws/bedrock/agentcore/prod-"
  - "/aws/bedrock/agentcore/staging-"
```

— Config key, file, and runnable YAML example all given. 5.41.0

03 **Ten new AWS checks including Bedrock AgentCore and GuardDuty**

NEW 80

Adds `iam_policy_passrole_to_bedrock_agentcore_restricted` (flags PassRole over roles reaching Bedrock AgentCore), `iam_policy_no_agentcore_workload_access_token_wildcard` (flags overly broad workload access token grants), `cloudwatch_log_group_agentcore_data_protection_policy_enabled` (verifies AgentCore log groups mask sensitive data), `guardduty_runtime_monitoring_enabled` (unified Runtime Monitoring across EC2/ECS/EKS), `guardduty_ai_protection_enabled` (AI Protection for Bedrock, AgentCore, SageMaker), `ecr_registry_enhanced_scanning_enabled` (Inspector-based continuous rescanning), `eks_cluster_vpc_cni_network_policy_enforced` (VPC CNI network policy enforcement), `iam_role_service_trust_restricts_source_to_account` (source-account constraint on service trust), `elasticbeanstalk_environment_no_secrets_in_configuration` (secret scanning in EB config), and `memorydb_cluster_in_transit_encryption_enabled` (TLS for MemoryDB).

— All ten check IDs and what each flags are named. 5.41.0

04 **Fifteen new checks across seven providers** NEW 80

Adds `ecr_repository_image_no_secrets` (scans latest ECR image config/filesystem for secrets), `bedrock_guardrail_contextual_grounding_filter_enabled`, `bedrock_custom_model_encrypted_with_cmk`, `bedrock_knowledge_base_encrypted_with_cmk`, `bedrock_agent_role_not_shared_across_agents`, `rolesanywhere_profile_restricts_session_permissions` (all AWS); `iam_workload_identity_pool_provider_attribute_condition` (GCP); `organization_default_workflow_permissions_read_only`, `repository_default_workflow_permissions_read_only`, `organization_actions_pull_request_approval_disabled` (GitHub); `defender_domain_dmarc_records_published` (M365 Defender, requiring `p=quarantine` or `p=reject`); `oss_bucket_versioning_enabled`, `oss_bucket_server_side_encryption_enabled` (Alibaba Cloud OSS); `vpc_security_group_open_egress` (Huawei Cloud); and `ske_cluster_no_public_endpoint` (STACKIT SKE, checking ACL and `0.0.0.0/0` / `::/0` CIDRs).

— All 15 check IDs and providers named with conditions. product docs

05 **Jira finding deep links via DJANGO_UI_BASE_URL** NEW

75

Adds `DJANGO_UI_BASE_URL` API environment variable; when set, Jira issues created from findings include a direct link back to the finding filtered by its UID.

Enable finding-to-issue deep links in Jira by pointing the API at your Prowler Cloud UI base URL.

```
$ DJANGO_UI_BASE_URL=https://cloud.prowler.com
```

— Named env var with a concrete setting example. 5.41.0

06 **Import Findings workflow for OCSF reports** NEW

75

Adds a browser-based Import Findings dialog on the Scans page (Prowler Cloud and Private Cloud) accepting `.ocsf.json` reports via drag-and-drop or file picker, with inline job-completion and error reporting (requires Manage Ingestions permission). Providers whose findings were loaded via this workflow now show an Imported provider indicator in the providers table, distinguishing them from live connected providers.

— UI path and file format named, permission requirement given. 5.41.0

07 **NCSC Cyber Essentials v3.3 compliance framework** NEW

75

Adds NCSC Cyber Essentials: Requirements for IT Infrastructure v3.3 (April 2026) as a universal compliance framework with 28 requirements across five control themes — Firewalls, Secure Configuration, Security Update Management, User Access Control, and Malware Protection — of which 16 requirements map to Azure checks and 12 are reported as Manual.

— Precise scope and numbers given, no direct usage step. product docs

08 **Insecure registry scanning via Trivy** NEW

70

Adds `--registry-insecure` flag propagated to Trivy via `TRIVY_INSECURE`, enabling pull and scan of images behind self-signed certificates on trusted internal registries.

— Exact flag and env var named, minimal added mechanism. 5.41.0

09 **Automation labels on Jira issues from findings** NEW

70

Jira issues created from findings now carry labels `prowler`, `prowler-<provider>`, `prowler-<severity>`, `prowler-<check-id>`, and `prowler-finding-<finding-uid>` for JQL filtering and automation matching.

— All label formats named but no usage walkthrough. 5.41.0

10 **MCP tool execution error reporting** IMPROVED 70

MCP server tools now report failures as proper MCP tool execution errors (`isError: true`) instead of successful responses carrying an `{"error": ...}` payload; affects `prowler_docs_search`, `prowler_docs_get_document`, `prowler_hub_get_check_code`, `prowler_hub_get_check_fixer`, and `prowler_get_compliance_framework_state_details`.

— Exact behavior change and affected tools named, no example call. 5.41.0

11 **Image provider registry scanning improvements** IMPROVED 65

Image provider registry scans now enumerate repositories in parallel, skip non-image OCI artifacts (Helm charts, cosign signatures, SBOM attestations), and continue scanning remaining images when Trivy fails on a single image. Authentication now falls back to Basic when a registry rejects a bearer token, and upgrades to bearer when the server responds with a Bearer challenge to a Basic or anonymous request.

— Clear before/after behavior, but no config surface to act on. 5.41.0

12 **Slack as an alert destination** NEW 65

Adds Slack as an alert destination: connect a workspace from the Integrations page, authorize destination channels, and select them in the alert modal's Destination channels selector alongside the existing Recipients email selector.

.

— Exact UI navigation given, no API surface named. product docs

THINNER COVERAGE BELOW

13 **CIS Google Workspace Foundations Benchmark v1.4.0** NEW 50

Ships the CIS Google Workspace Foundations Benchmark v1.4.0 compliance framework with reworked checks covering Gmail spoofing, 2-Step Verification, password expiration, and alert severity defaults.

— Names framework version and covered areas, no check IDs. 5.41.0

14 **Slack connection status shown on integration card** IMPROVED 50

Slack integration connection check now displays its result (including the refused or reached channels) directly on the integration card, not only in a transient toast notification (Prowler Cloud only).

— Describes the change but only a UI display tweak. 5.41.0

15 **Thumbs up/down feedback on Lighthouse AI answers** NEW 50

Adds per-answer thumbs-up/thumbs-down feedback with an optional description field to every Lighthouse AI chat response.

— Simple description, no field name or endpoint given. product docs

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Automox

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗

Check out our full list of release notes, and stay up to date with what

> Automox improved Exclusion Windows with local device time zone scheduling and more precise per-device conflict warnings to reduce false positives.

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

01 Local time zone scheduling for Exclusion Windows NEW

75

The Exclusion Window builder now includes a toggle to switch between Local time of device and UTC scheduling, so a window set for 2:00 AM fires at 2:00 AM in each endpoint's own time zone without manual UTC conversion. An interactive time preview in the builder shows the equivalent wall-clock time when toggling between Local and UTC, based on the current browser time zone.

Set a maintenance exclusion window to run at 2:00 AM in each device's local time zone, so globally distributed endpoints all get their local off-hours window without calculating per-region UTC offsets.

> 📍 In the Automox console, go to Manage › Exclusion Windows › New Exclusion Window, set your start time to 2:00 AM, then toggle the time zone selector to Local time of device before saving.

– Names exact toggle and UI path, includes usage steps

Local Exclusion Window Scheduling & Tar…

THINNER COVERAGE BELOW

02 Per-device conflict warnings for Exclusion Windows IMPROVED

45

Exclusion window warning modals now show per-device conflict counts (e.g., '3 of 12 selected devices are currently in an exclusion window') instead of group-level alerts, eliminating false-positive warnings.

– Describes behavior change but no exact UI path or config

Local Exclusion Window Scheduling & Tar…

WAS THIS USEFUL? 👍 👎

Fleet

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

Open device management

> Fleet introduces a configurable throttle for its own automated activities, preventing mass automation triggers from overwhelming hosts, while ensuring user-initiated actions always jump the queue.

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

01 Configurable release rate for Fleet-initiated activities NEW

83

A new `FLEET_ACTIVITY_FLEET_INITIATED_RELEASE_PER_MINUTE` environment variable throttles Fleet-initiated activities (policy-automation software installs, scripts, iOS/iPadOS scheduled app updates) to a configurable rate, defaulting to 1000 hosts/minute (0 = unlimited), to spread install-execution and result-ingestion load when automations fire across many hosts simultaneously. User-initiated activities (self-service installs, admin-run scripts, lock/unlock/wipe, setup experience) now take precedence over queued Fleet-initiated activities on the same host.

— Names exact env var, default value, and precedence behaviour. fleet-v4.90.2

WAS THIS USEFUL? 👍 👎

CrowdStrike

Feedback ⓘ 1 RELEASE · 2026-09-01 PRESS ↗

CrowdStrike is a cybersecurity platform providing cloud-based endpoint protection, threat intelligence, and incident response services.

> CrowdStrike launched SafeMind, an agentic offensive/defensive AI system native to the Falcon platform, built on NVIDIA Nemotron models and CoreWeave infrastructure.

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

01 **SafeMind agentic AI system in Falcon** NEW 68

CrowdStrike introduced SafeMind, an agentic system operating natively in the Falcon platform that pairs an offensive model, Red Tempest (trained for advanced attack-scenario emulation on CrowdStrike sensor telemetry, threat intelligence, Falcon Complete MDR event annotations, and fifteen years of incident response data), with a defensive model, Blue Solano (trained on the same dataset to autonomously deploy defensive measures), inside a closed-loop harness that continuously pits them against each other to improve. The harnesses also interoperate with third-party frontier and open-source models for model choice and cost control; SafeMind is built on NVIDIA Nemotron open models, trained and served on CoreWeave's AI Cloud, and will be made available to external parties as standalone models and harnesses through the Project QuiltWorks program. CrowdStrike reports SafeMind achieves a 29% higher detection rate, 6x faster end-to-end remediation, and 99% cost reduction versus leading frontier and open-source model baselines.

— Names models, infra and metrics but no API, flag or config to act on.

2026-09-01-898a0e58

WAS THIS USEFUL? 👍 👎

Trend Vision One

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

Trend Vision One is a unified cybersecurity platform providing threat detection, incident response, and risk management across endpoints, networks, and clouds.

> Trend Vision One's Container Security added new scan types and finer-grained policy controls for Runtime Malware Scanning, including image filtering and path exclusions.

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

01 **Filesystem path exclusions in Container Security policy** NEW

65

The Container Security policy adds configurable filesystem path exclusions for scanning; `/proc`, `/dev`, and `/sys` are excluded by default, and additional paths can be added to fine-tune which paths within scanned images are examined.

— Names default excluded paths verbatim, lacks config key or command

New malware scanning controls in Contai…

02 **Image filter rules in Container Security policy** NEW

60

The Container Security policy gains image filter rules that let you exclude specific images (e.g. trusted third-party operators) from scanning or restrict scanning to only matching images, with per-namespace scoping so the same image can follow different policies in different namespaces.

— Names mechanism and scoping but no exact policy field or UI path

New malware scanning controls in Contai…

THINNER COVERAGE BELOW

03 **Scheduled and real-time scan types for Runtime Malware Scanning**

NEW

45

Container Security Runtime Malware Scanning now supports two scan types: scheduled scans, aimed at comprehensive coverage against the latest malware patterns, and real-time scans, aimed at immediate threat detection right after deployment.

— Describes purpose of each scan type but no config path given

New malware scanning controls in Contai…

WAS THIS USEFUL? 👍 👎

Sophos Anti-Virus for macOS

Feedback ⓘ

1 RELEASE · SEEN 2026-09-02 NOTES ↗

Sophos Anti-Virus for macOS is endpoint security software that detects and removes malware, viruses, and threats on Apple computers.

> Sophos Anti-Virus for macOS 2026.2.0 adds support for macOS 27 Golden Gate and bumps the threat detection engine to v3.99.0.

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

01 **Support for macOS 27 Golden Gate** NEW 30

Sophos Anti-Virus for macOS now supports macOS 27 Golden Gate.

— Names the OS version but no detail on scope or changes Version 2026.2.0

02 **Threat detection engine updated to v3.99.0** IMPROVED 30

The threat detection engine has been updated to version 3.99.0.

— Names version number but no detail on what changed Version 2026.2.0

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Yeti

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

Your Everyday Threat Intelligence

> Yeti 2.11.0 brings configurable model choice and finer session control to the agent chat interface.

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

01 **Per-message model selection in agent chat** NEW 75

Agent chat now supports choosing a model per message, with the selection remembered per conversation so each session shows which model answered it. The list of selectable models is configured via `GEMINI_MODELS` in `prod/.env.example` (yeti-docker) or `agents.geminiModels` in the Helm chart.

— Names config keys and behavior but no exact UI steps 2.11.0

THINNER COVERAGE BELOW

02 **Agent chat session deletion** NEW 35

Adds the ability to delete agent chat sessions, removing the conversation and its metadata together.

— States what changed but no interface or endpoint detail 2.11.0

WAS THIS USEFUL? 👍 👎

IntelOwl

Feedback ⓘ 1 RELEASE · 2026-08-31 NOTES ↗

IntelOwl: manage your Threat Intelligence at scale

> IntelOwl v6.8.0 standardizes CTI connectors with a new base class, tightens up chatbot query handling and adds an Ollama keep-alive setting, and improves connector health-check error reporting in the UI.

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

01 Chatbot query handling and Ollama keep-alive IMPROVED

80

The chatbot now defaults `analyze_observable` to a curated playbook and enforces rejection of placeholder names to prevent accidental stub submissions, hardens jobs-vs-investigations routing, and accepts null filter arguments to broaden valid queries. A new `OLLAMA_KEEP_ALIVE` environment variable (e.g. `OLLAMA_KEEP_ALIVE=30m` in `.env`) keeps the Ollama model loaded in VRAM between chatbot requests to avoid cold-start delays.

Keep the Ollama model loaded in VRAM between chatbot requests to avoid cold-start delays in high-frequency analysis workflows.

```
$ OLLAMA_KEEP_ALIVE=30m
```

— Names env var, config file and multiple behavior changes. v6.8.0

THINNER COVERAGE BELOW

02 Connector health-check error messages IMPROVED

50

Connector health-check methods now return a (status, message) tuple, and the UI surfaces the full error message when a health check fails.

— Names return type change but no UI path given. v6.8.0

03 CTIConnector base class for connectors NEW

35

Introduces a `CTIConnector` base class to standardize CTI connector implementations and support data-model enrichment across them.

— Names the class but not how connectors use it. v6.8.0

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-09-02 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI's 7.260901.0 release adds cross-instance workflow status synchronization, built-in Filigran UI themes, and a new vulnerability data module.

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

01 **Workflow status sync via streams** NEW 35

Workflow release status can now be synchronized by name across instances through streams, enabling cross-instance status propagation without manual mapping.

– Names the mechanism (streams) but no config or command given. 7.260901.0

02 **Built-in Filigran themes** NEW 25

OpenCTI now ships built-in Filigran themes, giving teams out-of-the-box branded UI options without needing custom theme configuration.

– Describes benefit but no theme names or settings path given. 7.260901.0

03 **Vulnerability module in data model** NEW 25

A new vulnerability module has been added to the data model, expanding structured tracking of vulnerability entities.

– States the addition but no fields or schema details given. 7.260901.0

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Edgescan

Feedback ⓘ 3 RELEASES · SEEN 2026-08-10 → 2026-09-01

NOTES ↗

Edgescan is a vulnerability management and penetration testing platform that identifies and tracks security risks in applications and infrastructure.

> Edgescan's biggest move this window is Atomic, a new end-to-end automated penetration-testing engine that plans its own attacks and steadily expands coverage across injection, IDOR, business-logic, path-traversal and authenticated testing while surfacing live visibility into running assessments; alongside it, the core platform gained new EASM filters, per-location-type default scan profiles, TOTP-authenticated DAST scanning, and asset/organization management improvements.

└──▷ WHAT SHIPPED · 22 FEATURES most completely described first ? what's the number?

01 **Assessment run visibility: live stream, phase indicator, activity timeline**

IMPROVED 65

Adds a live assessment stream so findings and activity can be watched as an assessment runs (v0.1.0); adds a phase indicator so assessments show which phase is currently running, with access-control findings categorised more accurately and results saved even if an assessment ends early (v0.4.0); and adds a live assessment activity timeline giving a readable, real-time view of tester activity while an assessment is in progress (v0.5.0).

– Traces mechanism across three increments, no exact UI path

Release Notes - Atomic (pre-release)

02 **Broken access-control (IDOR) testing** IMPROVED 65

Adds broken-access-control (IDOR) modules (v0.1.0), then extends access-control testing to more link and request types and enables testing the same target as both a logged-in and anonymous user to expose access-control gaps (v0.4.0).

– Describes concrete before/after testing scope Release Notes - Atomic (pre-release)

03 **Business-logic testing for discount, coupon and quantity abuse**

IMPROVED 65

Adds a business-logic testing agent team covering abuse of discount and coupon handling in commerce flows (v0.2.0), adds detection for quantity manipulation such as

negative values to surface invalid order states (v0.3.0), and expands business-logic detection in commerce flows building on the discount/coupon and quantity-manipulation checks (v0.4.0).

— Names three concrete checks across increments Release Notes - Atomic (pre-release)

04 **EASM filtering by service, port, API and certificate** NEW

65

Adds a new EASM/ASM filter to scope the host and domain list to assets where no services were discovered, enabling investigation without a full CSV export; adds new EASM domain page filters to narrow domains by open port, API presence, and API confidence for precise external attack surface triage; and adds a certificate filter to filter certificates by issue type.

— Names every filter field but no exact UI path Release Notes - 25th Aug 2026 Release Notes - 10th Aug 2026

05 **Per-location-type default scan profiles** NEW 65

Adds support for two separate organisation-default scan profiles based on asset location type — one for URL-only unauthenticated web app assets and one for URL+Network — so new assets automatically inherit the correct profile.

— Names the two profile types and inheritance behaviour Release Notes - 25th Aug 2026

06 **Contact assignment at asset creation** IMPROVED 60

Adds ability to assign asset owner, technical POC, and assessment notifier emails when creating a single asset, eliminating the need to configure them post-creation.

— Names fields and describes before/after workflow Release Notes - 10th Aug 2026

THINNER COVERAGE BELOW

07 **Bulk role permission assignment** NEW 55

Bulk Actions on the organization index page now includes a create role permission option, enabling bulk permission assignment within an organization's scope.

— Names the UI feature and scope precisely Release Notes - 10th Aug 2026

08 **Injection detection coverage (XSS, SQL injection)** IMPROVED

50

Adds injection modules covering cross-site scripting, SQL injection, and similar (v0.1.0), then expands this detection with more complete and durably recorded results (v0.4.0).

— Names vulnerability classes across two increments Release Notes - Atomic (pre-release)

09 **Discovery phase using robots.txt and sitemaps** IMPROVED

50

Expands the discovery phase to use `robots.txt` and sitemap entries to find additional testable areas.

— Names concrete discovery sources Release Notes - Atomic (pre-release)

10 **Independent finding confirmation step** NEW 50

Introduces a separate, independent confirmation step for each potential finding before it is marked validated; findings that do not pass remain marked as suspected.

— Explains validation states but no config surface Release Notes - Atomic (pre-release)

11 **Multi-identity authentication and session persistence** NEW 50

Adds foundational multi-identity authenticated testing where an assessment can carry multiple sets of credentials as named identities, and strengthens session persistence through multi-step and session-rotating application flows to broaden coverage of authenticated areas.

— Two related auth mechanisms described together Release Notes - Atomic (pre-release)

12 **Path traversal and file disclosure testing** NEW 50

Adds path-traversal detection where requests are manipulated to test whether files outside the intended directory can be read, and adds broader file-disclosure testing support.

— Describes mechanism for one check, second is bare Release Notes - Atomic (pre-release)

13 **TOTP support in authenticated DAST scans** NEW 50

Adds TOTP support in authenticated DAST scans, enabling scanning of applications protected by MFA/time-based one-time passwords.

— Names the mechanism but no setup steps Release Notes - 25th Aug 2026

14 **Credentials Vault integration for authenticated assessments** NEW 45

Adds Credentials Vault integration so credentials can be passed to the harness for authenticated assessment of login-protected application areas.

— Names Credentials Vault but no config or endpoint detail

Release Notes - Atomic (pre-release)

15 **Retesting of previously reported findings** NEW 40

Adds a retesting capability that re-checks previously reported issues and reports whether they appear fixed or still present.

— Clear behaviour but no trigger or command named Release Notes - Atomic (pre-release)

16 **Pause and resume assessments** NEW 40

Adds pause and resume capability so assessments can safely stop and restart around maintenance windows or deployments.

– States use case but no mechanism Release Notes - Atomic (pre-release)

17 **EASM export content and clarity improvements** IMPROVED 40

Domain names are now included in EASM export files for records and services, and a new tooltip clarifies the difference between the index page export option and the bulk actions export option.

– Names export fields but minor scope Release Notes - 10th Aug 2026

18 **End-to-end automated pentest assessments** NEW 35

Adds end-to-end automated assessments where the tester builds its own attack plan and produces findings without per-step human direction.

– Describes capability but no mechanism or named surface

Release Notes - Atomic (pre-release)

19 **WAF-adaptive testing** NEW 35

Adds WAF-adaptive testing where assessment behaviour adapts based on how the application and any web application firewall respond.

– Explains behaviour but no scope or mechanism detail

Release Notes - Atomic (pre-release)

20 **Finding de-duplication and weakness mapping** NEW 35

Adds de-duplication and mapping to standard weakness categories for consistent triage.

– Bare description of a triage aid Release Notes - Atomic (pre-release)

21 **Custom request headers support** NEW 30

Adds support for custom request headers in assessments.

– Names a config-like capability with no further detail

Release Notes - Atomic (pre-release)

22 **Assessment export support** NEW 20

Adds export support for completed assessments.

– No format, endpoint or scope given Release Notes - Atomic (pre-release)

WAS THIS USEFUL? 👍 👎

Cloudmersive

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗

Cloudmersive provides cloud-based APIs for document processing, data validation, virus scanning, and content conversion tasks.

> Cloudmersive's Cloud DLP AI API now supports scanning much larger files submitted via multipart form uploads.

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

01 **Large file support for DLP scanning** IMPROVED 45

The Cloud DLP AI API now supports `multipart/form-data` encoded file uploads over 2 GB in size for DLP scanning.

— Names format and size limit but no endpoint or usage detail.

Cloudmersive Cloud DLP AI API v2-9-0 Re…

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ 2 RELEASES · 2026-09-01 → 2026-09-02 CHANGELOG ↗

runZero is a network asset discovery and inventory tool that identifies and maps all devices connected to networks in real-time.

> runZero's biggest changes this window were alert rules that can trigger on Issue lifecycle events and a new `comment:` search operator for the Issues inventory, alongside smaller additions spanning MCP OAuth, CBOM export, integration scripting, and reporting.

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

01 `comment:` search operator for Issues NEW

65

The `comment:` search operator is now supported in the Issues inventory for searching Issue comments, e.g. `comment:CVE-2025` or, via the search bar, `comment:<term>` to find Issues containing a specific term in their comments.

Search Issue comments for a keyword to quickly surface relevant discussion on open Issues.

```
$ comment:CVE-2025
```

Find Issues that contain a specific term in their comments — useful for tracking discussion-driven triage without leaving the Issues inventory.

📍 In the Issues inventory search bar, enter: comment:<term>

– Exact search syntax and runnable examples given runZero 5.1.260901.0 (2026-09-01) changelog-20260902-e664697f

THINNER COVERAGE BELOW

02 Alert rules for Issue lifecycle events NEW

50

Alert rules now support conditions for Issue lifecycle events, enabling automated responses to Issue state changes.

– Names condition type but no specific states or setup steps

runZero 5.1.260901.0 (2026-09-01) changelog-20260902-e664697f

03 External Reporting sharing for standard reports NEW

40

Standard reports can now be shared via the External Reporting feature.

— States the capability but no mechanism or steps `runZero 5.1.260901.0 (2026-09-01)`

04 **Faster scan processing via MAC merging optimization** IMPROVED

30

Large scan jobs now process faster after an optimization for MAC merging.

— Names the mechanism but gives no numbers or steps `changelog-20260902-e664697f`

05 **Censys Platform v3 integration** NEW

20

runZero added support for Censys Platform v3.

— Names the integration and version, no setup detail `product docs`

06 **MCP OAuth support** NEW

15

runZero added OAuth support for MCP integrations.

— Only named in summary, no detail on setup or scope `changelog-20260902-e664697f`

07 **WMI scripting support** NEW

15

runZero added WMI scripting capability.

— Named only, no mechanism or usage detail `changelog-20260902-e664697f`

08 **Mermaid charting support** NEW

15

runZero added Mermaid diagram charting support.

— Named only, no context on where charting appears `changelog-20260902-e664697f`

09 **AI-based labelling** NEW

10

runZero added AI-based labelling capability.

— Only named in summary, no scope or mechanism given `product docs`

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Aembit

Feedback ⓘ 1 RELEASE · SEEN 2026-09-02 PRESS ↗

Aembit is IAM for agentic AI and workloads. Enable secretless access and real-time policy enforcement across any environment to protect data and accelerate AI adoption.

> Aembit shipped a new integration with CrowdStrike Falcon AI Detection and Response (AIDR) that lets MCP traffic for AI agents be inspected and governed using centrally configured policies enriched with Aembit workload identity.

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

01 **CrowdStrike Falcon AIDR integration for MCP traffic** NEW

70

Aembit's MCP Identity Gateway now applies centrally pre-configured CrowdStrike Falcon AI Detection and Response (AIDR) policies across all MCP clients, inspecting MCP tool listings, inputs, and outputs for AI agent traffic without requiring per-client configuration or custom agent-side inspection logic. Configurable policies support allow, report, block, or transform actions based on identity context and risks identified by AIDR. Aembit supplies run-time identities for both user-delegated and autonomous agents to AIDR to enrich its policy decisions, and CrowdStrike AIDR findings now link directly to corresponding Aembit access records so investigators can pivot between platforms without manually reconciling logs.

— Describes mechanism and scope well but no config keys or commands given.

launch-20260902-d7ab4bbe

WAS THIS USEFUL? 👍 👎

Didit

Feedback

1 RELEASE · 2026-09-02

CHANGELOG ↗

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit's biggest window addition is a full transaction-monitoring rules Management API alongside opt-in biometric-template retention with a breaking change to session-deletion response bodies; it also shipped new fraud signals, age-assurance templates, marketplace access, an upgraded Copilot, and in-console support tickets.

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

01 Opt-in biometric-template retention after deletion NEW

95

Retention of biometric templates after session deletion can be enabled via `face_retention_policy` and `face_retention_days` on `PATCH /v3/webhook/`, and overridden per call with `retain_face_embeddings` on `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/`. Retained templates carry a finite `expires_at` and never contain images or identity fields.

Enable biometric-template retention for an app so duplicate detection and Face Search keep working after session data is purged.

```
$ curl -X PATCH https://api.didit.me/v3/webhook/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"face_retention_policy": "retain", "face_retention_days":
365}'
```

Delete a session while issuing a privacy-erasure instruction to purge all retained biometric templates for that person, regardless of the app-level retention policy.

```
$ curl -X DELETE
https://api.didit.me/v3/session/<sessionId>/delete/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"deletion_instruction": "privacy_erasure",
"instruction_id": "erasure-req-20260902"}'
```

— Exact field names plus runnable curl examples for both endpoints

changelog-20260902-2e1362ed

02 **Transaction-monitoring rules Management API** NEW 80

New `/v3/transactions/rules/` Management API endpoints let you create, update, delete, and backtest transaction-monitoring rules programmatically; matching `library` and `install` endpoints manage preset rules.

– Names full endpoint path but no example call given changelog-20260902-2e1362ed

03 **Session-deletion endpoints now return 200 with JSON body** 75

BREAKING

`DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body (`face_retention_outcome` , `biometric_template_uuid` ; per-session `results[]` for batch) instead of `204 No Content` — clients that assert on `204` must be updated.

– Clear before/after contract change but no migration example changelog-20260902-2e1362ed

04 **Didit Copilot upgraded to floating panel** IMPROVED 65

Didit Copilot is now a floating panel available on every console page, with inline source citations, a fullscreen mobile layout, workflow-building from a description, Database Validation configuration by conversation, and questionnaire generation from a CSV or Excel file.

– Lists many capabilities but no exact command or config changelog-20260902-2e1362ed

THINNER COVERAGE BELOW

05 **Marketplace and per-feature engine selection** NEW 55

The Marketplace is now accessible in the Business Console, with per-feature engine selection available in both the graph and simple workflow editors.

– Names UI area but no config or endpoint detail changelog-20260902-2e1362ed

06 **ISO age-assurance workflow templates** NEW 40

Ready-made age-assurance templates are now available in the workflow builder, with returned data already restricted.

– Only names the feature, no mechanism or config detail changelog-20260902-2e1362ed

07 **Manual company entry in verification flows** IMPROVED 40

Manual company entry was added to web, iOS, and Android verification flows for use when a registry provider is unavailable.

– Thin description, no config or endpoint named changelog-20260902-2e1362ed

└──▷ **BREAKING ON UPGRADE**

! `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body (`face_retention_outcome` , `biometric_template_uuid` ; per-session `results[]` for batch) instead of `204 No Content` — clients that assert on `204` must be updated.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-09-02 CHANGELOG ↗

The easiest way to add authentication and user management to your application. Purpose-built for React, Next.js, Remix, and “The Modern Web”.

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

01 Reverification window setting for sensitive actions NEW

75

Adds a Reverification window setting (1–10 minutes, default 10) under Session lifetime on the Sessions page in the Clerk Dashboard, controlling how long a successful reverification remains valid for Clerk-protected sensitive actions like account deletion.

Reduce the reverification window so users must re-verify more frequently before sensitive actions like deleting their account.

> 📍 In the Clerk Dashboard, go to Sessions › Session lifetime › Reverification window and set the value to a number between 1 and 10 minutes.

— Exact setting, range, default, and dashboard path given. changelog-20260902-634d332a

02 Biometric sign-in APIs for mobile SDKs NEW

75

Adds biometric credential APIs across Clerk's mobile SDKs: `useBiometricCredentials` for React Native/Expo, `Clerk.shared.biometricCredentials.signInWithBiometrics` and `Clerk.shared.biometricCredentials.biometricCredentialId` for Swift, and `Clerk.biometricCredentials.signInWithBiometrics` and `Clerk.biometricCredentials.biometricCredentialId` for Android. These enable enrollment and sign-in via Face ID, Touch ID, or Android biometrics using a device-bound private key.

— Named APIs per platform but no usage example given. changelog-20260902-634d332a

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗

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

> Auth0 has made Advanced MFA Configurations generally available, letting admins customize OTP lifetimes for SMS and email-based multi-factor authentication.

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

01 **Configurable OTP lifetimes for MFA** NEW 60

Advanced MFA Configurations are now Generally Available, enabling customizable OTP lengths, lifetimes, and device-remember windows. Specifically, OTP lifetime for SMS MFA can be set from 30 seconds to 1 hour, and for email MFA from 5 minutes to 1 hour.

.
.
.
.

— Names concrete lifetime ranges but no config path or API given.

Advanced MFA Configurations are now Gen…

WAS THIS USEFUL? 👍 👎

Okta for AI Agents

Feedback ⓘ 1 RELEASE · 2026-09-02

CHANGELOG ↗

Okta for AI Agents provides identity and access management capabilities enabling secure authentication and authorization for artificial intelligence agents.

> Okta for AI Agents added JA4 TLS fingerprinting across Syslog events and manual MCP server registration for legacy servers, alongside a new suspicious authenticator enrollment detection, an added OAuth scope, new OIN integrations, and a Jamf Pro SCIM mapping field.

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

01 JA4 TLS fingerprinting in Syslog events NEW

88

Okta now captures JA4 TLS client fingerprints in `securityContext.tlsFingerprint.ja4` across all Syslog event types, including telephony (OTP/SMS), sign-in, auth, and token events, enabling detection of bot traffic, toll fraud, and TLS-based attack patterns beyond IP/user-agent signals alone. Not yet available for orgs on custom-hosted domains.

Query the System Log to surface JA4 TLS fingerprints across all event types and identify bot or toll-fraud patterns that IP/user-agent signals miss.

```
$ curl -s -H 'Authorization: SSWS <api_token>' \
  'https://<okta_domain>/api/v1/logs?
filter=securityContext.tlsFingerprint.ja4+pr+"JA4:..."&limit=100'
```

— Names exact field and includes a runnable System Log query. changelog-20260902-6b98254f

02 Manual MCP server registration NEW

83

Adds manual MCP server registration, allowing admins to configure authorization server details and client credentials directly for internal or legacy MCP servers that don't support automated metadata discovery endpoints. In the Admin Console, go to AI Agents > MCP Servers > Add MCP Server, select Manual registration, then enter the authorization server URL, client ID, and client secret.

Register an internal MCP server that lacks automated metadata discovery by supplying authorization server details and client credentials manually.

> 📍 In the Admin Console, go to AI Agents › MCP Servers › Add MCP Server, select Manual registration, then enter the authorization server URL, client ID, and client secret for the internal or legacy MCP server.

— Exact UI navigation and required fields given. changelog-20260902-6b98254f

03 **Application username format field for Jamf Pro** NEW

65

Adds Application username format field to the Jamf Pro integration in the Admin Console by default, enabling custom mappings for the SCIM `userName` attribute.

— Names exact field and attribute but no example values. changelog-20260902-6b98254f

THINNER COVERAGE BELOW

04 **AI agent binding tab shows authentication requirement**

IMPROVED

50

Displays the authentication requirement for an app on the AI agent binding tab, with a direct link to the app's Sign On tab to configure an authentication policy.

— Names the UI location as a clear starting point. changelog-20260902-6b98254f

05 **Suspicious authenticator enrollment detection** NEW

45

Adds Suspicious authenticator enrollment detection to Okta Identity Threat Protection (ITP), flagging suspicious activity during authenticator enrollment with details recorded in the System Log for investigation.

— Names the detection and log location but no steps or thresholds.

changelog-20260902-6b98254f

06 **New OIN integrations added** NEW

45

Adds new OIN integrations: Flexopus (SAML), Paminga (OIDC), Paminga (SCIM), Upwave Lifecycle Management by Redblock (SCIM), Velantix Axiom (OIDC), and Velantix Axiom (SCIM).

— Lists all named integrations but no configuration detail. changelog-20260902-6b98254f

07 **interclient_access scope for authorization servers** NEW

35

Adds `interclient_access` scope to the scope list for authorization servers.

— Names the scope but gives no usage detail. changelog-20260902-6b98254f

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Netskope

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

Netskope is a cloud security platform that provides visibility and threat protection for cloud applications and web traffic.

> Netskope R141 introduces per-app-instance visibility controls for SSPM findings, letting teams scope findings to individual SaaS app deployments.

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

01 **Per-app-instance visibility for SSPM findings** NEW

33

Adds the ability to control visibility of SSPM findings scoped to individual app instances, enabling teams to restrict or expose findings by specific SaaS app deployments.

— Describes capability but no config key, endpoint, or UI path given.

Netskope Cloud 141.0.0 Documentation Av…

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

ScoutDNS

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

DNS filtering and threat protection for businesses, MSPs, schools, and nonprofits. Block malware, phishing, and unwanted content at the DNS layer.

> ScoutDNS's Windows Client 2.2.6 adds compatibility with Microsoft Entra Global Secure Access, letting ScoutDNS and Entra Private Access coexist without manual configuration.

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

01 **Microsoft Entra Global Secure Access compatibility** NEW

70

Windows Client 2.2.6 (driver 1.0.0.22) adds compatibility with Microsoft Entra Global Secure Access (GSA), allowing ScoutDNS and Entra Private Access to coexist without disrupting private DNS or tunnel startup. It automatically detects an active GSA Private Access tunnel and permits only required GSA bootstrap and synthetic DNS paths while continuing to enforce ScoutDNS policy for public DNS, with no manual domain or IP exclusions required. Validated in both split tunnel (GSA Private Access) and full tunnel modes.

— Names driver version and tunnel modes but no config steps

Windows Client 2.2.6 adds Microsoft Ent…

WAS THIS USEFUL? 👍 👎

Knocknoc

Feedback ⓘ 1 RELEASE · 2026-09-02 CHANGELOG ↗

Knocknoc is a security tool for detecting and managing security vulnerabilities in software dependencies and code.

> Knocknoc's 26.08 release focused on expanding dynamic network access control to new backend targets, adding native integrations for Fastly, nftables, and Kemp LoadMaster.

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

01 **New backend integrations for Fastly, nftables, and Kemp LoadMaster**

NEW 65

Knocknoc added three new backend integrations for dynamic, just-in-time network access control: a Fastly integration that controls which source IPs are permitted through a Fastly edge service (also described as an ACL integration for network access control enforcement), enabling JIT access at the CDN layer; a native `nftables` integration that dynamically adds and removes source IPs via nftables rules on Linux hosts when users authenticate; and a Kemp LoadMaster integration that dynamically controls network access through the load balancer based on Knocknoc authentication events.

— Names each backend and its mechanism, but no config keys or commands given

changelog-20260902-ac6830d6

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

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

> CISO Assistant's v3.21.4 release adds automatic SSO user provisioning and broadens workflow tooling and localization.

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

01 **Just-in-Time SSO user provisioning** NEW 50

User accounts are now automatically created on first SSO login, eliminating the need for manual pre-provisioning of accounts before a user can sign in via SSO.

— Explains mechanism but no config key or setup steps given v3.21.4

02 **Expanded workflow enablers and sample workflows** IMPROVED 20

Adds more workflow enablers and additional sample workflows to the platform's workflow capabilities.

— Vague description with no named enablers or workflows v3.21.4

03 **Updated Dutch translations** IMPROVED 20

Adds updated Dutch (nl) translations for the UI.

— Names the language but no scope of coverage v3.21.4

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

GitHub CLI

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗

GitHub's official command line tool

> GitHub CLI adds a repeatable `--attach` flag for uploading local images and videos directly into issues, PRs, and comments, alongside worktree support for `gh issue develop` and agent-friendly spinner output.

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

01 **Attach local images and videos via CLI** NEW 93

Adds a repeatable `--attach` flag to `gh issue create`, `gh issue edit`, `gh issue comment`, `gh pr create`, `gh pr edit`, and `gh pr comment` that uploads a local image or video to GitHub and embeds the resulting URL inline in the body; the same file cannot be attached twice. Alt text for images is supported via a `#` suffix on the file path (e.g. `--attach 'PATH/TO/IMAGE#ALT-TEXT'`), though not for videos. When combined with `--body` or `--body-file` (or stdin/interactive editor input), GitHub CLI rewrites any matching local file references already present in the Markdown body in-place with the uploaded URL rather than appending a duplicate, and appends new ones otherwise. Requires push access to the repository.

Add a screenshot inline to a PR when opening it, so reviewers see the image rendered without leaving the CLI.

```
$ gh pr create \
  --title "Fix login screen layout" \
  --body-file ./pr-description.md \
  --attach ./screenshots/before.png \
  --attach ./screenshots/after.png
```

Attach an image with descriptive alt text to an issue comment for accessibility.

```
$ gh issue comment 42 --attach 'screenshots/error.png#sign-in
screen showing authentication error'
```

Embed a video player in an issue comment by referencing the local file in the body and attaching it, so GitHub CLI rewrites the path to a rendered player.

```
$ gh issue comment 42 \
  --body-file ./repro-video-body.md \
  --attach ./recordings/repro.mp4
```

— Full flag syntax, behavior, and constraints documented with runnable examples. v2.99.0

02 **Worktree support in gh issue develop** NEW 70

Extends `gh issue develop` with a `--worktree` flag, used alongside `--checkout`, to create a linked branch and check it out in a new Git worktree without disturbing the current working copy.

Create a linked feature branch for an issue and check it out in an isolated worktree, leaving your main checkout untouched.

```
$ gh issue develop 123 --checkout --worktree /path/to/wt-feature
```

— Names exact flags and gives a runnable example command. v2.99.0

THINNER COVERAGE BELOW

03 **Agent-readable spinner output** IMPROVED 23

GitHub CLI now outputs text-only spinner output when `gh` is invoked by a coding agent, improving agent-readable output.

— Only a brief prose description with no mechanism or usage detail. v2.99.0

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 2, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue