THE CYBER TOOLCHAIN

NO. 041

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 10, 2026 · EVERY WEEKDAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 23 tools

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

Nmap ⓘ changes since 2026-07-11 CODE ↗

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

Nmap gains TLS close_notify tolerance, SSL service portrule coverage, and signal propagation to child processes with `-k`.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-11
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-11
```

└──▷ TRY IT

Specify an interface with `-e` to bind both the outgoing interface and the source address, useful when a host has multiple NICs and you need traffic to leave a specific one.

```
$ nmap -e eth1 -S 192.168.1.50 <target>
```

- Adds `-k` signal propagation: termination signals (INT/TERM/HUP) are now caught and forwarded to the entire process group, then Ncat exits with `128 + signum`, ensuring child processes are not orphaned.
- NSE scripts now include `s` (SSL-protected) service names — such as `amt-soap-https` — in their portrules, so scripts run against SSL-wrapped services as well as their standard unencrypted counterparts.
- Nmap now tolerates TLS peers that do not send `close_notify`, preventing spurious scan failures against non-compliant TLS targets.
- `http-referer-checker` now categorizes but does not fetch off-host URLs, limiting requests to URLs on the target host only.

WAS THIS USEFUL? 👍 👎

web-check

1 RELEASE · 2026-07-28 NOTES ↗

All-in-one OSINT tool for analysing any website

> Web-check 2.2.0 adds new security-focused panels to its scan results, a subdomain lookup check, and the ability to disable individual checks via configuration.

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

01 **Vulnerabilities and Advisory panels in results UI** NEW

37

A new Vulnerabilities panel with improved WAF detection surfaces web application firewall findings for scanned targets, alongside a new Advisory panel that displays security advisories directly in the results UI.

— Names two panels and WAF detection but no mechanism detail 2.2.0

02 **Configurable check disabling** NEW

28

Adds configuration support for disabling individual checks, letting users opt out of specific scans.

— States the capability but no config key or file named 2.2.0

03 **Subdomain lookup check** NEW

26

Adds subdomain lookup as a new check capability for scanned websites.

— Bare description of a new check, no config or output details 2.2.0

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

adscan

3 RELEASES · 2026-07-16 → 2026-08-03 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.

> Across three releases ADscan added end-to-end AD CS ESC5 exploitation, integrated password cracking backed by a 94-million-entry wordlist, new SQL Server and Shadow Credentials attack paths, resumable scans with native macOS/ARM64 support, and a batch of new commands (`harvest` , `jobs` , `generate_report` , `writeup` , `execute attack_paths`) — alongside breaking changes to the REPL command set and a move to container-only execution.

└──▷ WHAT SHIPPED · 33 FEATURES most completely described first ? what's the number?

01 **execute attack_paths subcommand** NEW 85

New `adscan execute attack_paths` subcommand runs a single saved attack path or step against a workspace without a full rescan; resolves credentials automatically for already-owned accounts, and a companion re-arm resets step status to re-run a path.

Re-run a single saved attack path against a workspace after remediating one step, without repeating the full scan.

```
$ adscan execute attack_paths
```

— Named subcommand with runnable example and re-run mechanism v10.1.0

02 **Integrated password cracking with wordlists and effort ladder**

NEW 85

Adds password-strength auditing built into the scan: credentials captured via broadcast poisoning, Kerberos service tickets, and AS-REP hashes are cracked in the background using the bundled 94-million-entry 'ADscan AD Audit Wordlist' and a client-targeted wordlist generator, with recovered passwords fed back into the attack path automatically. A four-rung selectable cracking-effort ladder is calibrated to host hardware with accurate pre-commit time estimates and live progress/ETA, and is configurable via the YAML scan-config file and web platform.

Run a fully unattended CI scan with cracking effort and poisoning capture configured ahead of time via the YAML scan-config file.

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

— Concrete wordlist size and effort tiers, config file named generically v10.0.0

03 **writeup command for lab/CTF reports** NEW 75

New `writeup` command captures port data, directory contents, the full attack chain as an editable Mermaid diagram, timestamped steps with credential provenance, flags, and dead ends into a lab/CTF writeup; runs automatically on lab and CTF scans.

— Command named with detailed contents, auto-trigger behaviour noted v11.0.0

04 **Client-logo branding for reports** NEW 75

Adds `adscan deliver --client-logo` and `adscan ci --client-logo` flags accepting PNG, SVG, or JPG to embed a client logo in both the free and PRO report covers.

Generate a client-branded exposure report during a non-interactive CI scan, embedding the client's logo in both the HTML and PDF output.

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

— Exact flags and a runnable example provided v11.0.0

05 **SQL Server attack coverage** NEW 75

Adds SQL Server attack coverage: linked-server lateral movement across domain and forest trusts, automated command execution with escalation to SYSTEM, and arbitrary file read via bulk import — each surfaced as a validated attack step in the report; the SQL Server post-authentication workflow now runs automatically in unattended scans instead of defaulting to skip.

— Multiple named techniques, no direct CLI flag shown v10.1.0

06 **ESC5 end-to-end AD CS exploitation** NEW 70

Executes AD CS ESC5 end to end: backs up the CA private key, forges a certificate offline, and authenticates with it from the CA host's machine account or a recovered local-admin password.

— Full mechanism described but no runnable command shown v11.0.0

07 **generate_report command for exposure reports** NEW 70

New `generate_report` command produces a free self-contained HTML and paginated A4 PDF exposure report from any scan; audit-mode scans run it automatically.

— Named command and output formats given, no example used `v11.0.0`

---

08 **Telemetry and offline controls** NEW 70

Adds `--no-telemetry` and `--offline` flags to `adscan start`; also restores `set telemetry off` as a per-workspace setting and introduces `set telemetry off global` for an operator-wide default.

— All flags and commands named verbatim `v11.0.0`

---

09 **Attack path collapsing and refined result states** IMPROVED 70

Attack paths with interchangeable intermediate accounts now collapse into a single finding that names and counts those accounts, letting the operator choose which to target — several at once for a roast or spray, one for a direct step — with unattended runs picking the strongest candidate automatically; attack paths also carry more precise result states: 'Partially Validated', 'Not Executed for Safety', and 'Attack Surface Reduced — Hardening Observed'.

— Mechanism and exact state names given, no direct command `v11.0.0` `v10.1.0`

---

10 **REPL flag-form arguments and tab-completion** IMPROVED 70

REPL commands now accept flag-form arguments `-d` / `--dc-ip` / `-u` / `-p` / `--host` / `--service` alongside the existing positional form, with tab-completion for both flag names and values.

Use flag-form arguments in the REPL for scripted or tab-completed input instead of positional syntax.

```
$ start_auth -d 10.0.0.1 -u alice -p 'S3cr3t!' --host corp.local
```

— Exact flags named with a runnable REPL example `v10.1.0`

---

11 **harvest command for recovered credentials** NEW 70

Adds a `harvest` command providing a single view of every credential obtained during a scan (captured, cracked, and sprayed), each tagged with its privilege tier and reachable assets, plus an end-of-scan harvest summary.

Pull up every credential recovered so far — captured hashes, cracked passwords, and sprayed accounts — with privilege tier and reach, mid-engagement.

```
$ harvest
```

— Named command with runnable example and clear scope `v10.0.0`

12 **Broadcast poisoning capture** NEW 70

Adds broadcast name-resolution poisoning capture (LLMNR, NBT-NS, mDNS) as a background task, surfaced as a finding and linked attack step, with an explicit consent panel in interactive mode and automatic start in unattended scans; poisoning controls are configurable via the YAML scan-config file and web platform.

— Protocols and behaviour named, no direct trigger command `v10.0.0`

13 **jobs command for background tasks** NEW 65

Adds a `jobs` command to list, review, and stop long-running background tasks such as broadcast-name-resolution poisoning capture and password cracking.

Review and stop background cracking or poisoning jobs without interrupting the active scan.

```
$ jobs
```

— Named command with runnable example `v10.0.0`

14 **Writable-share credential bait** NEW 60

Adds opt-in writable-share credential bait: plants a bait file on writable shares and captures resulting authentication attempts as a background job that does not block the scan; configurable via new scan-configuration toggles in the CLI and web console.

— Opt-in mechanism and config surfaces named generically `v10.1.0`

THINNER COVERAGE BELOW

15 **COVERAGE.md technique-to-MITRE catalog** NEW 55

New `COVERAGE.md` lists all 103 Active Directory techniques implemented, each mapped to MITRE ATT&CK and marked executed, detected, or declined for safety; generated from the product catalog so it cannot drift.

— File name and scope given, no usage example `v11.0.0`

16 **Shadow Credentials escalation** NEW 55

Adds Shadow Credentials escalation: detects when it can write a key credential to a target account and executes the full authentication chain end to end.

— Mechanism described, no command or config named `v10.1.0`

17 **NTLM credential capture via pivot** NEW 55

Adds credential capture through a pivot: NTLM authentication can be captured across a routed tunnel with per-target selection of the callback address.

— Mechanism named, no command or flag given v10.1.0

18 **Native Certificate Services attack implementation** IMPROVED 55

Certificate Services attacks now run through a native implementation, including detection of delegated CA-management rights and detection of Extended Protection on the web-enrollment endpoint (suppressing false ESC8 findings when enrollment is hardened).

— Mechanism and false-positive fix described, no command v10.1.0

19 **Built-in engine replaces bundled external tools** BREAKING 55

Password spraying, share enumeration, SMB/LDAP hardening checks, local credential verification, and account cleanup now run on ADscan's built-in engine with consistent proxy support, removing the dependency on bundled external command-line tools; any workflow that relied on the previous bundled third-party tools directly must now use these native engine equivalents.

— Scope of replaced functions listed, no migration command given v10.0.0

20 **Cross-scan exposure comparison view** NEW 50

The Enterprise platform adds a cross-scan exposure comparison view that plots exposure across a workspace's scans and lists findings and attack paths that appeared, closed, or carried over.

— UI capability named, no exact navigation path given v11.0.0

21 **Container-only execution enforced** BREAKING 50

Running the ADscan core directly on the host machine is now refused; the tool runs only inside its container.

— Clear before/after behaviour, no command specifics v10.1.0

22 **Extended-rights escalation chaining** IMPROVED 50

Extended-rights escalation (forced password reset, DCSync, LAPS password read) is now modeled per target and chained into attack paths.

— Named rights listed but no execution detail v10.1.0

23 **ConvertFrom-SecureString secret recovery** NEW 50

Adds recovery of secrets protected with PowerShell `ConvertFrom-SecureString -Key`, chained into share-password findings.

— Names exact cmdlet but no usage flow shown v10.1.0

24 **Native-Microsoft-tooling remediation guidance** IMPROVED 50

Remediation guidance is now written entirely in native Microsoft tooling (`dsacls`, RSAT Active Directory module, `Get-*` / `Set-*` cmdlets) with no third-party product names in client-facing report text.

— Names exact tools/cmdlets, no example command run v10.1.0

25 **Resumable scans** NEW 50

Adds resumable scans — interrupted scans (crash, disconnect, or deliberate stop) can resume from the last completed phase or host; the web platform can stop and resume a running scan gracefully.

— Behaviour explained, no command or flag given v10.0.0

26 **GPO control chained to code execution** NEW 45

GPO control is now chained through to code execution via a scheduled task on machines the GPO applies to.

— Mechanism named briefly, no configuration detail v10.1.0

27 **Privilege Blast Radius reporting** NEW 45

Adds Privilege Blast Radius: reports, the web platform, and the CLI now show the full set of assets and capabilities a single compromised account reaches.

— Concept described, no surface or command named v10.0.0

28 **Native macOS and ARM64 support** NEW 45

Adds native macOS and Apple Silicon support; LITE and PRO images are now published for ARM64/Apple Silicon in addition to x86-64.

— Platforms and image tiers named, no install command v10.0.0

29 **HTTP SPN relay and coercion detection** NEW 40

Service accounts carrying an HTTP SPN are now flagged as a Kerberos web-application relay and coercion surface, with native remediation guidance in the report and platform.

— Short description, no configuration or command detail v11.0.0

30 **Expanded web port sweep** IMPROVED 30

Web ports 8080 and 8443 are now included in the important-ports sweep and folded into service discovery.

— Exact ports named but change is minor and automatic v10.1.0

31 **Refined AD exposure score** IMPROVED 30

The AD exposure score now combines proof strength with per-technique exploitability.

— One-line description with no scoring mechanics v10.1.0

32 **REPL commands renamed to native names** BREAKING 25

Nine REPL commands were renamed to native names; the previous command names are no longer accepted.

— Count given but no old/new names listed v10.1.0

33 **'Blocked by Active Controls' report category removed** BREAKING 25

The 'Blocked by Active Controls' report category was removed from reports.

— Single-line removal notice with no migration detail v10.1.0

└──▷ **BREAKING ON UPGRADE**

! Nine REPL commands were renamed to native names; the previous command names are no longer accepted.

! The 'Blocked by Active Controls' report category was removed.

! Running the ADscan core directly on the host machine is now refused; the tool runs only inside its container.

! Bundled external third-party command-line cyber tools have been removed; any workflow that relied on them directly must now use ADscan's native engine equivalents.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

npq

2 RELEASES · 2026-08-09 → 2026-08-10 NOTES ↗ CODE ↗

safely install npm packages by auditing them pre-install stage

> npq added machine-readable audit output for both scripted and agent-driven workflows, via an explicit `--json` flag and automatic detection of coding-agent environments.

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

01 **`--json` flag for audit-only output** NEW 75

The new `--json` CLI flag emits a versioned JSON audit report and never proceeds to install, letting CI and scripted pipelines consume npq results programmatically. It routes invocations to an audit-only path that bypasses interactive prompts and auto-continue behaviour, and the JSON report captures incomplete audit check results so callers can distinguish fully-evaluated from partially-evaluated signals.

Audit a package in CI and capture the machine-readable result without triggering an install — useful for policy checks or SIEM ingestion.

```
$ npq install express --json
```

— Named flag with runnable example, but no schema details given. v3.24.0

THINNER COVERAGE BELOW

02 **Auto-detection of coding-agent environments** NEW 55

npq now detects when it is running inside a coding-agent environment and automatically defaults both audit output and package parsing to JSON mode, without requiring manual flags. In this mode, package parsers switch to a JSON-safe mode that rejects inputs like credential-bearing URLs that are unsafe to process as JSON.

— Explains mechanism and safety behaviour but no manual trigger or flag given. v3.25.0

WAS THIS USEFUL? 👍 👎

Vercel deepsec

changes since 2026-07-11

CODE ↗

Deepsec is a security harness for finding vulnerabilities in your codebase powered by coding agents

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-11 https://github.com/vercel-labs/deepsec.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-11
```

- Ships one-shot Deepsec initialization (`coordinator.ts`) to scaffold a project from scratch without manual setup steps.
- Introduces stable, deterministic `findingId` values (SHA-256 of projectId + normalized path + title) assigned at append time and lazily backfilled on old records, so finding identity survives across revalidation runs.
- Adds permissive verdict reconciliation for `revalidate` : verdicts now match by alias/findingId → exact title → normalized title → unique 1:1 remainder, with a `matchedBy` diagnostic recorded per match, replacing the previous exact-equality-only approach that silently dropped mismatched verdicts.
- Adds short per-invocation prompt aliases (`F1` … `Fn`) mapping to stable finding IDs so agents never see raw hashes in prompts.
- Introduces attack-surface coverage analysis (`coverage.ts`) with typed `SurfaceKind` (`http` , `rpc` , `queue` , `cron` , `cli` , `webhook` , `agent-tool` , `other`) and `SurfaceExposure` (`public` , `authenticated` , `internal` , `mixed` , `unknown`) inventories, enabling structured surface-aware scanning.

+1 more

WAS THIS USEFUL? 👍 👎

DependencyCheck

1 RELEASE · 2026-08-03 NOTES ↗

OWASP dependency-check is a software composition analysis utility that detects publicly disclosed vulnerabilities in application dependencies.

> DependencyCheck v13.0.0 adds OpenSSL 3.x detection while removing Nexus V2 Search support and previously deprecated configuration options.

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

01 **OpenSSL 3.x detection in OpenSSLAnalyzer** NEW 40

The `OpenSSLAnalyzer` now supports detection of OpenSSL 3.x, enabling accurate identification of modern OpenSSL versions in scanned dependencies.

— Names the analyzer but no detection mechanism detail. v13.0.0

02 **Removal of Nexus V2 Search and deprecated config** BREAKING 35

Nexus V2 Search support has been dropped along with all previously deprecated configuration keys and integration options. Any working setup relying on Nexus V2 or these deprecated settings will break on upgrade.

— States impact but not which specific config keys removed. v13.0.0

└──▷ BREAKING ON UPGRADE

! Nexus V2 Search support has been removed along with all previously deprecated configuration keys and integration options — any working setup relying on Nexus V2 or deprecated settings will break on upgrade.

WAS THIS USEFUL? 👍 👎

sbomify

1 RELEASE · 2026-07-30 NOTES ↗

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

> sbomify shipped its first public REST API (187 endpoints across 17 areas) alongside a major VEX/CBOM expansion — full VEX authoring, triage and drift detection, CycloneDX 1.7 CBOM ingestion, and new Trust Center posture visibility.

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

01 **Public REST API with 187 endpoints** NEW 78

sbomify now publishes a full REST API — 187 endpoints across 17 areas, documented at `https://app.sbomify.com/api/v1/openapi.json`. Named areas include Artifacts (23 endpoints, create/read/update/delete), Controls (21 endpoints, create/read/update/delete), Cle (20 endpoints, create/read), Workspaces (20 endpoints — manage workspace settings, members, branding, and collaboration features), Products (19 endpoints — structure software inventory into products grouped by releases), Releases (19 endpoints — tag and manage product releases with downloadable artifacts and version tracking), Compliance (17 endpoints, create/read/update), Components (11 endpoints — organize software components that contain SBOMs and documents), plus 9 more areas: Vulnerability Scanning, Access Requests, Oidc Trusted Publishing, Plugins, Billing, Licensing, Notifications, User, and Internal.

— Names all 17 areas, counts and spec URL, but no per-endpoint paths 26.7.1

02 **End-to-end VEX authoring, triage, and drift detection** NEW 78

Adds a full VEX (Vulnerability Exploitability eXchange) workflow: authoring, uploading, and ingesting VEX documents in CycloneDX, OpenVEX, and CSAF formats; in-product VEX triage and lifecycle management with an opt-in Dependency-Track triage sync that auto-pins to releases; and scheduled drift detection that flags stale VEX statements and uncovered High/Critical findings.

— Full mechanism and formats named, but no config keys or commands v26.7.1

03 **Version and format filters on SBOM list endpoint** IMPROVED 68

The component SBOMs list endpoint gains `version` and `format` filtering parameters, letting API consumers narrow returned SBOM listings by these fields.

— Names exact query parameters on a specific endpoint v26.7.1

04 **CBOM ingestion with CycloneDX 1.7 normalization** NEW

65

Adds CBOM (Cryptography Bill of Materials) ingestion across every lineage, with CycloneDX 1.7 registry normalization and a new workspace dashboard for cryptographic asset visibility.

— Names format version and dashboard but no exact UI path v26.7.1

THINNER COVERAGE BELOW

05 **VEX and CBOM downloads in Trust Center** IMPROVED 50

Trust Center now exposes VEX and CBOM downloads alongside a new vulnerability posture panel, giving external viewers visibility into exploitability and cryptographic bill-of-materials data.

— Clear feature location in Trust Center, limited mechanism detail v26.7.1

06 **Workspace-scoped token binding on Workspaces API** BREAKING 48

Enforces workspace-scoped token binding across the workspaces API, restricting API tokens to the workspace they were issued for.

— Names affected API surface but no migration guidance given v26.7.1

07 **Per-artifact usage analytics via PostHog** NEW 40

Adds PostHog per-artifact analytics tracking VEX, CBOM, and HBOM uploads, downloads, deletions, and release pulls.

— Names tool and tracked events, no dashboard or access path v26.7.1

08 **OpenGrep SAST added to CI pipeline** NEW 30

Adds an OpenGrep SAST workflow to sbomify's own CI pipeline.

— Bare mention of internal tooling change, no user-facing detail v26.7.1

WAS THIS USEFUL? 👍 👎

cve-lite-cli

2 RELEASES · 2026-07-14 → 2026-07-26 NOTES ↗

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

> cve-lite-cli added dependency-maintenance detection (`--check-maintenance` /DM001), an offline advisory database bootstrap, release-cooldown-aware fix warnings, and SARIF output for GitHub Code Scanning, alongside three new override-hygiene rules and refinements to the overrides subcommand and HTML report.

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

01 --check-maintenance dependency maintenance checks NEW

91

New `--check-maintenance` flag (rule `DM001`) flags direct dependencies that block a transitive CVE fix via major-version constraint drag, or that are deprecated on npm. Findings render in terminal, JSON, multi-folder, and HTML report output, and count toward `--fail-on` (high severity for constraint drag, medium for deprecated packages).

Catch direct dependencies that are blocking a transitive CVE fix or are npm-deprecated, and fail CI on any high-severity finding.

```
$ cve-lite /path/to/project --check-maintenance --fail-on high
```

— Names flag, rule id, severities and every output surface. v1.28.0

02 SARIF output for GitHub Code Scanning NEW

86

New `--sarif` output includes `ruleIndex`, `kind: "open"`, and `partialFingerprints` fields for full GitHub Code Scanning compatibility, with shared fingerprint logic extracted to `src/utils/sarif.ts`.

Generate a SARIF report with full GitHub Code Scanning fingerprinting to upload as a code-scanning artifact in CI.

```
$ cve-lite /path/to/project --sarif
```

— Named flag, exact SARIF fields, source file, and runnable command. v1.27.0

03 Cooldown-aware fix version warnings NEW

85

New fix resolver reads `min-release-age` (npm `.npmrc`), `minimumReleaseAge` (pnpm), and `npmMinimalAgeGate` (yarn) and warns when a recommended fix version is newer than the configured release window. It is advisory only, on by default, does not change flags or exit codes, and renders in terminal, `--verbose`, JSON, and HTML report for npm/pnpm/yarn (Bun not yet covered).

— Names every config key and package manager but has no example command. v1.28.0

---

04 **New override hygiene rules PD001, PD002, OA009** NEW 83

Rule `PD001` (Override-only phantom, high) fires when source code imports a package that is only present via an override pin, prompting declaration as a direct dependency. Rule `PD002` (Transitive-only phantom, medium) fires when source code imports a package that is only present as a transitive dependency, prompting an explicit declaration. Rule `OA009` (Stale floor detector) fires when a `>=` or `^` override floor is already met by every parent declaration, flagging the override as redundant, with a safety guard that suppresses the stale-floor removal suggestion when the override anchors a phantom import (PD001/PD002 interaction). A new `examples/all-scenarios/` fixture triggers all nine auto-detectable override hygiene rules in a single scan.

— Names three rule ids with firing conditions and a fixture, no direct command. v1.27.0

---

05 **overrides subcommand JSON output and path-first ordering**

IMPROVED 70

The `overrides` subcommand now supports `overrides --json` output written to a timestamped file, consistent with the main scan path, and accepts path-first argument ordering so `cve-lite <path> overrides [flags]` works alongside the existing `cve-lite overrides <path> [flags]` form.

Run override hygiene auditing with auto-fix on a non-root project path using the new path-first argument order.

```
$ cve-lite ./packages/api overrides --fix
```

Export override hygiene findings to a timestamped JSON file for ingestion into a SIEM or audit log.

```
$ cve-lite /path/to/project overrides --json
```

— Two named CLI surfaces with runnable example commands for each. v1.27.0

---

06 **`advisories init` offline advisory database** NEW 65

New `advisories init` subcommand creates an empty offline advisory database for air-gapped environments, to be populated manually and used with runtime-API-free scanning (e.g. paired with `--offline`).

Bootstrap an empty offline advisory database on an air-gapped machine, then scan without any runtime API calls.

```
$ cve-lite advisories init && cve-lite /path/to/project --offline
```

— Names the exact subcommand and a runnable pairing, but mechanism is thin. v1.28.0

THINNER COVERAGE BELOW

07 **HTML report skipped-dependency and theme improvements** IMPROVED

43

The HTML report and `report.json` now show skipped manifest dependencies, and the HTML report gained a light/dark mode toggle persisted to `localStorage`.

— Two thin report tweaks named but no usage path given. v1.28.0 v1.27.0

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2 ⓘ changes since 2026-07-20 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds Remus malware detection, a threat-intelligence processing module (ThreatFox + Malpedia), and PCI vendor ID spoofing to defeat Al-Khaser VM checks.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-20
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-20
```

- New `modules/processing/threatintelligence.py` processing module (runs at `order=22`) queries ThreatFox for contacted IPs/domains and Malpedia for malware family cards, annotating network hosts with threat tags and writing results to `results['threatintelligence']` with `providers`, `indicators`, `families`, and `stats` keys.
- New Remus malware family detection and dynamic config extraction added to CAPE's detection pipeline.
- New PCI vendor ID spoofing capability to bypass Al-Khaser QEMU detection checks (`Qemu.cpp` vendor ID fingerprint).
- New `require_task_delete` and `can_delete_task` authorization guards added to the multitenancy access-control layer, gating task deletion to submitter or box-admin roles.
- Adds per-task `try/except` in APIv2 batch delete loops so partial-apply failures are reported with a `partial_error` field rather than returning a bodiless 500.

+1 more

└──▷ BREAKING ON UPGRADE

! The `suspended` parameter has been removed from execute() methods in `process.py` and `abstracts.py`; any custom modules or scripts passing `suspended=` will break on upgrade.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

kubescape

1 RELEASE · 2026-07-22 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 v4.0.11 adds a GitLab SAST report format, a quarantine remediation action, CEL-based admission control scanning, an AWS ECR vulnerability adaptor, and expands patch export targets, alongside several MCP server and policy-scoping enhancements.

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

01 **GitLab SAST output format** NEW 70

Adds `gitlab-sast` as a new `--format` output option for Kubescape's printer, enabling native integration with the GitLab Security Dashboard.

— Names exact flag/value but no example output shown v4.0.11

02 **CEL-based admission control scanning** NEW 70

`kubescape scan` now runs CEL-based admission controls as part of the scan, bundling and evaluating the VAP CEL policy suite without requiring a network call.

— Explains mechanism and offline behavior, no exact command v4.0.11

03 **VAP library defaults to embedded bundle** BREAKING 60

The `vap deploy-library` subcommand now serves the embedded VAP bundle by default; network download only occurs when `--from-release` is passed, a behavior change from the previous default of downloading from the network.

— Names subcommand, flag, and exact before/after behavior v4.0.11

THINNER COVERAGE BELOW

04 **Quarantine remediation action** NEW 55

Adds a `quarantine` action to the `kubescape operator remediate` CLI subcommand for isolating non-compliant workloads.

— Names subcommand and action but not full syntax v4.0.11

05 **AWS ECR vulnerability adaptor** NEW 55

Implements an AWS ECR vulnerability adaptor via the new Container Image Vulnerability (CIV) Adaptor interface, enabling ECR-native image vulnerability data ingestion.

— Names interface and provider but no usage steps v4.0.11

06 **Expression-based exception scoping** NEW 55

Honors `objectSelector.matchExpressions` in `PostureExceptionPolicy.ObjectSelector`, enabling expression-based exception scoping.

— Names exact field path, no worked example v4.0.11

07 **OCI and local export for patch command** IMPROVED 50

`kubescape patch` now supports OCI and local export targets, expanding patched image delivery options beyond remote registries.

— Names command and targets but no flag syntax given v4.0.11

08 **MCP server scanning and profiling additions** NEW 50

The MCP server gains a headless OPAProcessor for low-latency Network Policy scanning, a headless OPAProcessor RBAC scanner, and new container profile tools and resources.

— Names three additions but no invocation details v4.0.11

09 **Encrypted resource metadata decryption** NEW 35

Adds decryption support for encrypted resource metadata (report crypto), complementing the existing encryption path.

— Only prose description, no config or command named v4.0.11

10 **Reversible container metadata anonymization** NEW 25

The anonymizer now supports reversible container metadata transformation.

— Bare description, no mechanism or usage given v4.0.11

└──▷ **BREAKING ON UPGRADE**

! The `vap deploy-library` subcommand no longer downloads the VAP bundle from the network by default — it serves the embedded bundle instead. Pass `--from-release` to restore the previous download behaviour.

WAS THIS USEFUL? 👍 👎

kata-containers

1 RELEASE · 2026-07-20 NOTES ↗

Kata Containers is an open source project and community working to build a standard implementation of lightweight Virtual Machines (VMs) that feel and perform like containers, but provide the workload isolation and security advantages of VMs. https://katacontainers.io/

> Kata Containers 4.0 rewrites the runtime in Rust (`runtime-rs`) as the new default, ships the Dragonball VMM built-in, and adds dm-verity/EROFS integrity and storage support alongside broad device, storage and networking enhancements across QEMU, Cloud Hypervisor and Dragonball.

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

01 **Rust runtime rewrite (runtime-rs) becomes default** BREAKING

72

Kata Containers 4.0 introduces `runtime-rs`, a full Rust rewrite of the Kata Containers runtime, now the default runtime replacing the original Go runtime, which is deprecated; `kata-deploy` now sets the RuntimeClass to `runtime-rs` by default and existing users may see minor differences in configuration and behavior when upgrading. `runtime-rs` now supports x86_64, ARM (aarch64) and IBM s390x, and the containerd shim is updated to version 0.11 with ttrpc updated to version 0.9.

— Names runtime, dependency versions and migration impact, no exact command `4.0.0`

02 **Block-based emptyDir volumes and sizeLimit enforcement** BREAKING

60

Adds block-based `emptyDir` volumes (`block-plain`) for improved performance and security by avoiding filesystem sharing. The previously unsafe `emptyDir` setting that allowed volumes to exceed their `sizeLimit` has been removed.

— Names volume type and removed unsafe sizeLimit behavior `4.0.0`

03 **Unified block device model and storage improvements** IMPROVED

60

Adds a unified block device model (BlockModern) across QEMU, Cloud Hypervisor, and Dragonball, with improved hot-plug/unplug reliability and rollback on failure. Also adds discard and unmap support for thin-provisioned storage (requires hypervisor support) and improved `virtio-scsi` support as a container root filesystem driver.

— Explains mechanism across hypervisors, no exact flags `4.0.0`

THINNER COVERAGE BELOW

04 **EROFS snapshotter support** NEW

55

Adds support for the EROFS snapshotter, available through Helm configuration, with `kata-deploy` validating EROFS prerequisites before installation.

— Names deployment path via Helm and kata-deploy validation 4.0.0

05 **VFIO device passthrough enhancements** BREAKING 55

Adds VFIO device cold-plug support, including CDI-based VFIO cold-plug in single-container (non-pod-sandbox) mode, and s390x VFIO-AP mediated device support for IBM Z crypto and accelerator workloads. Unused VFIO configuration options have been removed.

— Names modes and architectures, no config specifics 4.0.0

06 **Networking enhancements across hypervisors** IMPROVED 55

Propagates multi-queue networking configuration through all supported hypervisors (QEMU, Cloud Hypervisor, Dragonball), adds support for placing network devices in the host network namespace when using QEMU, and improves network interface hot-plug for QEMU with enhanced SR-IOV support.

— Names hypervisors and behavior, no config keys 4.0.0

07 **Dragonball VMM as built-in default** NEW 50

Adds Dragonball as a built-in VMM default, supporting `virtio-net`, `virtio-blk`, and VFIO device passthrough without requiring QEMU or Cloud Hypervisor.

— Names supported devices, no config path given 4.0.0

08 **dm-verity integrity-verified rootfs with GPT/VMDK support** NEW 45

Adds dm-verity integrity-verified root filesystem support, enabled by default in deployment configurations, plus GPT/VMDK support for integrity-verified root filesystems.

— Describes feature and formats but no config keys 4.0.0

09 **GPU device passthrough support** NEW 35

Adds GPU device support across QEMU, Cloud Hypervisor, and Dragonball hypervisors.

— Names hypervisors only, no further detail 4.0.0

10 **Improved static sandbox resource sizing** IMPROVED 35

Improves static sandbox resource sizing to account for runtime overhead values, benefiting Kubernetes Pod Overhead and scheduling accuracy.

— Explains benefit, lacks mechanism specifics 4.0.0

11 **virtio-mem memory hotplug in runtime-rs** NEW 30

Adds `virtio-mem` memory hotplug support in `runtime-rs`.

— Bare feature name with no further detail `4.0.0`

└──▷ **BREAKING ON UPGRADE**

! The original Go runtime is deprecated and `kata-deploy` sets the RuntimeClass to `runtime-rs` by default — existing users may encounter minor differences in configuration and behavior when upgrading.

! Unsafe `emptyDir` setting that allowed volumes to exceed their `sizeLimit` has been removed.

! Unused VFIO configuration options were removed.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

Tenzir ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Tenzir is a security data platform that ingests, transforms, and analyzes security telemetry and threat intelligence data.

> Tenzir v6.10.0 introduces a suite of native time-series analytics functions for beacon detection and threat-hunting cadence analysis, alongside deeper ClickHouse export support and automatic NetFlow/IPFIX ingestion.

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

01 Beacon-detection time-series analytics functions NEW

90

Adds a set of TQL functions for cadence and periodicity analysis aimed at beacon/C2 detection: deltas() computes successive differences of a sorted timestamp list into inter-arrival durations; mad() (available both as a `summarize` aggregation and as a list method) computes median absolute deviation, returning a duration when given duration input; skewness() computes moment skewness with Bowley skewness selectable via `method="bowley"` (returns `0.0` for zero-dispersion input); autocorrelation(xs, max_lag=int) computes normalized autocorrelation coefficients for lags 0 through `max_lag` (default half the list length); periodogram(xs) returns spectral power per period via FFT; dominant_period(times, resolution=duration) returns the strongest period and a normalized strength (0-1); and ewm_variance()/ewm_stddev() join ewma() with the same decay options plus a `bias` option to switch between debiased and biased weighted moments.

Identify periodic beacons by dominant period strength — filter destinations with a strong 30-second check-in pattern from noisy traffic.

```tql
summarize dst, times=collect(t), samples=count()
beat = times.dominant_period(resolution=5s)
where samples >= 8 and beat.strength >= 0.8
select dst, period=beat.period, strength=round(beat.strength *
100) / 100
```

Detect beacon periodicity in connection logs by computing the dominant period of inter-arrival deltas.

```
$ from connections | summarize times=collect(ts) by src_ip,
dest_ip | extend d=deltas(times) | extend
period=dominant_period(autocorrelation(d))
```

Identify irregular beaconing candidates by flagging sources whose inter-arrival skewness or MAD exceeds expected thresholds.

```
$ from connections | summarize times=collect(ts) by src_ip |
extend d=deltas(times) | where mad(d) > 5s or skewness(d) > 2.0
```

— Names every function signature and parameter with runnable examples

Tenzir Node v6.10.0: Native time-series …

02 **ClickHouse export and import enhancements** IMPROVED 73

`to_clickhouse`'s `json` and `low_cardinality` arguments now accept nested fields (e.g. `json=file.xattributes`) rather than only top-level ones, and gain automatic JSON column detection, including detection after `ocsf::cast`, so freeform OCSF fields are automatically typed as ClickHouse JSON columns. `from_clickhouse` can now read ClickHouse's native `JSON` column type, fixing a previous serialization version error.

— Names exact config keys and fixed error, no example given

Tenzir Node v6.10.0: Native time-series …

03 **Automatic NetFlow/IPFIX format detection** IMPROVED 60

`read_auto` now automatically detects and selects `read_netflow` for NetFlow v5, NetFlow v9, and IPFIX byte streams, removing the need to specify the format explicitly on ingestion.

Auto-ingest a NetFlow v9 capture file without specifying the format explicitly.

```tql
from_file "capture.nfv9" { read_auto }
```

— Names the function and formats with a runnable ingestion example

Tenzir Node v6.10.0: Native time-series …

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One

2 RELEASES · 2026-08-10 BLOG ↗

Trend Vision One is Trend Micro's unified cybersecurity platform providing threat intelligence, vulnerability management, and incident response across endpoints and networks.

> Trend Vision One added configurable cloud classification timing controls and AWS China support to File Security Storage, while the Zero Trust Secure Access on-premises gateway upgraded its prompt injection detection engine and expanded AI/MCP access logging.

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

01 Configurable classification timeouts and AWS China support NEW

90

File Security Storage template v4.4.0 adds the `FileSecurityStorageScannerICRCQueryTimeoutMs` deployment parameter to set a configurable timeout for cloud classification queries (default 6000 milliseconds), and the `FileSecurityStorageScannerICRCQueryRetryCount` parameter to set a configurable retry count for timed-out queries (default 3 retries; set to `0` to disable retries). It also adds support for deploying the template to AWS China (`aws-cn`) regions.

— Names exact deployment parameters, defaults, and region support

File Security Storage template version 4…

THINNER COVERAGE BELOW

02 Prompt Scan engine and AI/MCP log enrichment IMPROVED

45

Zero Trust Secure Access On-premises Gateway 1.0.90 adopts the latest Prompt Scan engine to improve detection accuracy for prompt injection and harmful prompts, and enriches activity and detection logs with greater visibility into AI service access and MCP server access.

— Describes engine upgrade and log enrichment without configuration detail

Zero Trust Secure Access On-premises Gat…

WAS THIS USEFUL? 👍 👎

Perplexity numbat

1 RELEASE · 2026-07-29 NOTES ↗

Visibility into AI agent activity on endpoints, with on-device detection, optional pre-action blocking, and forensic reconstruction.

> numbat's initial v0.1.1 release ships a complete CLI toolchain for AI agent endpoint visibility: hook-based live monitoring and OTLP collection, forensic scan/timeline reconstruction from on-disk artifacts, a CEL/YAML rules engine with opt-in pre-action blocking, portable case bundles, and a versioned NDJSON record schema.

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

01 **CEL and sequence rules engine** NEW 85

Adds `numbat rules check`, `numbat rules list`, and `numbat rules test --fixture FILE` for validating, compiling, and evaluating CEL and multi-step sequence rules against NDJSON event fixtures. Adds `--rules-dir DIR` (repeatable) and `--no-builtin-rules` flags to `scan`, `collect`, `hook install`, and `rules check|list|test` for loading operator-authored YAML rules or replacing built-in rules by id. Ships built-in CEL and multi-step sequence detection rules evaluated locally, alongside support for custom YAML rules in operator-controlled directories.

Enforce a custom operator policy that blocks SSH authorized_keys writes by Codex — copy the shipped rule, add `enforce: true`, validate, then reinstall with enforcement enabled.

```
$ numbat rules check --rules-dir ./numbat-policy && numbat hook install --agent codex --emit all --rules-dir ./numbat-policy --enforce
```

— Names every rules subcommand and shared flag v0.1.1

02 **Hook lifecycle management for live agent integrations** NEW 75

Adds `numbat hook install --agent NAME|all` and `numbat hook uninstall --agent NAME|all` to wire and remove live integrations for supported desktop, CLI, IDE, and gateway AI agents, plus `numbat hook status [--agent NAME|all]` to inspect numbat-owned live integration configuration.

Install live monitoring for Codex and stream all events, findings, indicators, and enforcement decisions to the default record file.

```
$ numbat hook install --agent codex --emit all
```

— Names exact install/uninstall/status commands and flags `v0.1.1`

---

03 **Forensic scan and timeline reconstruction from agent artifacts**

NEW

75

Adds `numbat scan [--agent NAME ... | --path FILE|DIR ...]` for read-only forensic reconstruction from on-disk agent session artifacts without prior instrumentation, and `numbat timeline [--agent NAME ... | --path FILE|DIR ...]` to reconstruct per-session chronological views from the same artifacts. Scanning supports secret redaction, with raw evidence files opt-in only via case bundles.

— Two commands with flags and redaction behaviour specified `v0.1.1`

---

04 **Opt-in pre-action enforcement** NEW

75

Adds opt-in pre-action blocking via `--enforce` on `hook install`; rules with `enforce: true` in an operator `--rules-dir` emit enforcement decisions, while all shipped built-in rules remain monitor-only by default.

Enforce a custom operator policy that blocks SSH authorized_keys writes by Codex — copy the shipped rule, add `enforce: true`, validate, then reinstall with enforcement enabled.

```
$ numbat rules check --rules-dir ./numbat-policy && numbat hook install --agent codex --emit all --rules-dir ./numbat-policy --enforce
```

— States flag, rule field, and default posture `v0.1.1`

---

05 **Portable case bundles with manifest verification** NEW

70

Adds `numbat case build <case-id> --from FILE` and `numbat case verify DIR` to create and verify portable case bundles with SHA-256 manifests.

— Exact subcommands and integrity mechanism named `v0.1.1`

---

06 **Versioned NDJSON record schema and emission** NEW

65

Emits versioned NDJSON records (schema v0.2.0) for events, findings, enforcement decisions, indicators, and scan summaries, with JSON Schemas published under `docs/schema/v0.2.0/`. Adds `--emit all` flag to `hook install` to write events, findings, indicators, and enforcement decisions to `~/.numbat/records.ndjson`.

— Names schema version, record types, and output path `v0.1.1`

---

07 **Live OTLP log collection** NEW

60

Adds `numbat collect [--addr HOST:PORT]` to receive OTLP/HTTP protobuf logs from supported agents and emit normalized records.

— Names command and flag but limited mechanism detail v0.1.1

08 **Local agent discovery** NEW 60

Adds `numbat agents [--all]` for read-only discovery of supported local agents, reporting config signals, artifacts, live capture mode, and wired status.

Discover which AI agents are present on the endpoint before installing any hooks.

```
$ numbat agents --all
```

— Lists exact fields the command reports v0.1.1

THINNER COVERAGE BELOW

09 **Record shipping to remote endpoint** NEW 55

Adds `numbat ship --input-file F --http-url U` to tail a local NDJSON record file to an HTTP endpoint for remote delivery.

— Exact flags given but mechanism thinly described v0.1.1

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Cloudmersive ⓘ 1 RELEASE · 2026-08-09 BLOG ↗

Cloudmersive is an API platform providing cloud-based document processing, image manipulation, and data conversion services for developers.

> Cloudmersive's Private Cloud CDR API gained two new sanitization options for stripping metadata and tracking pixels from documents.

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

01 **Metadata and tracking pixel removal in CDR API** NEW 25

The Content Disarm and Reconstruction (CDR) API now supports a metadata removal option and a tracking pixel removal option, expanding the sanitization choices available when processing documents.

— Names two options but no parameter, flag, or endpoint details given.

Cloudmersive Private Cloud CDR API v2-13…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

ConductorOne

1 RELEASE · 2026-08-10 CHANGELOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance workflows across enterprise systems.

> ConductorOne shipped a unified Identities & NHI inventory dashboard, TypeScript typing for the functions SDK, inline access context in review tasks, and reworked notification toggles.

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

01 **Reworked notification settings** IMPROVED 60

Notification settings were updated: the 'your access request was created' confirmation is now toggleable per channel (email, Slack, Teams) and on by default; provisioning notifications moved under a dedicated section; always-on system notifications are no longer shown as a toggle.

— Specific channels and default named, but no config key given `changelog-20260810-1a3518d1`

THINNER COVERAGE BELOW

02 **Inline resource access in task details** IMPROVED 45

The task details drawer now shows what the reviewed resource has access to inline, so reviewers no longer need to leave the task to gather context.

— Describes the change but no named field or endpoint `changelog-20260810-1a3518d1`

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Clerk

1 RELEASE · 2026-08-10 CHANGELOG ↗

Clerk is a user authentication and identity management platform that provides sign-up, login, and user profile management for web and mobile applications.

> Clerk shipped a beta Client ID Metadata Document (CIMD) mechanism for OAuth clients, Google Workspace Directory Sync, and out-of-the-box strict enumeration protection for combined sign-in-or-up flows.

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

01 **CIMD beta for OAuth clients** NEW 85

New CIMD (Client ID Metadata Document) beta lets OAuth clients use an HTTPS URL as their `client_id` — Clerk fetches and validates the metadata document, eliminating the need for pre-issued client IDs, secrets, or Dynamic Client Registration; aimed at MCP and other public OAuth clients. A new 'CIMD Clients' tab on the OAuth provider page lets admins explicitly allow clients by Client ID URL, assign scopes, review admission status and metadata fetch health, refresh metadata, and block or delete saved clients. An 'Advertise CIMD support' toggle publishes CIMD in authorization server metadata so any compatible client can connect, and pairs with 'Only allow pre-registered clients to connect' to restrict OAuth flows to pre-reviewed clients.

— Names client_id mechanism, UI tab, and toggles precisely changelog-20260810-8c78fd95

02 **Google Workspace Directory Sync** NEW 79

Directory Sync now supports Google Workspace: Clerk syncs users, groups, and group memberships via a Google service account with domain-wide delegation, without requiring sign-in events to trigger updates. It is configured in the Dashboard under a Google SAML connection's 'Directory sync' tab by uploading a service account key and specifying the admin to impersonate. Adds `organizations.title` and `customSchemas.<schema>.<field>` (e.g., `customSchemas.EmployeeInfo.costCenter`) as mappable custom attribute paths, including array paths for multi-valued attributes, and synced Google groups can be mapped to Clerk Roles.

— Names config paths, dashboard tab, and attribute keys changelog-20260810-8c78fd95

03 **Strict enumeration protection with sign-in-or-up** IMPROVED 76

`<SignIn />` now supports strict user enumeration protection and sign-in-or-up together out-of-the-box, with no code changes required — verification occurs before Clerk decides whether to sign in or create an account. Adds `signUpIfMissing` option for custom sign-in flows to enable sign-in-or-up behavior without enabling strict

enumeration protection globally, plus a `sign_up_if_missing_password_preferred` console warning on development instances when `<SignIn />` detects an incompatible strict-protection + password configuration. Strict enumeration protection is enabled on the Dashboard's Rules page under Protect; existing apps using `<SignIn />` as a combined sign-in-or-up page with public sign-ups continue working with no migration.

— Names option, warning key, and dashboard location `changelog-20260810-8c78fd95`

WAS THIS USEFUL? 👍 👎

Okta for AI Agents ⓘ 1 RELEASE · seen 2026-08-10 NOTES ↗

Okta for AI Agents provides identity and access management capabilities for securing AI agent applications and their interactions with enterprise systems.

> Okta for AI Agents entered beta with new discovery capabilities for spotting AI agents across Microsoft Copilot Studio and endpoints, plus an early-stage Agent Gateway component.

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

01 AI agent discovery for Copilot Studio and endpoints NEW

38

Identity Security Posture Management (ISPM) can now discover AI agents built with Microsoft Copilot Studio, and a new Endpoint AI agent discovery capability identifies AI agents running across endpoints.

— Names two discovery surfaces but no mechanism or setup steps given. `snapshot-20260810`

02 Agent Gateway beta capability NEW

21

Adds Agent Gateway as a new beta capability within the Okta for AI Agents platform, described only at a high level with no further mechanism given.

— Bare name only, no detail on what the gateway does. `snapshot-20260810`

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Firezone

Gateway 1.6.0 (2026-08-10) · 3 RELEASES · 2026-08-10 · BLOG ↗

LATEST · STABLE

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

Gateway 1.6.0 enables GRO/GSO on Linux 6.2+ TUN devices, roughly doubling single-stream TCP throughput through the tunnel.

- Enables GRO and GSO on the TUN device on Linux kernels 6.2 and later, roughly doubling single-stream TCP throughput through the tunnel.
- Lowers the system-wide `net.core.wmem_max` to 2 MB and `net.core.rmem_max` to 16 MB on the Gateway, reducing UDP send and receive socket buffer sizes.
- Speeds up recovery after all relays are lost by immediately requesting a fresh relay set from the portal.
- Improves reconnection speed after a network change.

└──▷ BREAKING ON UPGRADE

! The Gateway now sets `net.core.wmem_max` to 2 MB (down from a previously higher value) and `net.core.rmem_max` to 16 MB system-wide, which reduces UDP socket buffer sizes and may affect other processes sharing those kernel parameters on the same host.

▾ 2 more releases in this issue · 2026-08-10

Linux GUI 1.5.16 (2026-08-10) · BLOG ↗ · STABLE

Linux GUI 1.5.16 enables GRO/GSO on the TUN device for Linux 6.2+, roughly doubling single-stream TCP tunnel throughput.

- Enables GRO and GSO on the TUN device on Linux kernels 6.2 and later, roughly doubling single-stream TCP throughput through the tunnel.
- Speeds up recovery after all relays are lost by immediately requesting a fresh set from the portal.

Linux Headless 1.5.11 (2026-08-10) · BLOG ↗ · STABLE

Firezone Linux Headless 1.5.11 enables GRO/GSO on TUN devices for Linux 6.2+, roughly doubling single-stream TCP tunnel throughput.

- Enables GRO and GSO on the TUN device on Linux kernels 6.2 and later, roughly doubling single-stream TCP throughput through the tunnel.
- Speeds up recovery after all relays are lost by immediately requesting a fresh set from the portal.
- Improves reconnection speed after a network change.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

OpenCode ⓘ 1 RELEASE · 2026-08-10 NOTES ↗ CODE ↗

The open source coding agent.

> OpenCode v1.18.16 improves the desktop Home screen's project handling and makes config file parsing more forgiving of unknown fields.

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

01 **Project management in desktop Home screen** IMPROVED 40

Right-clicking a project in the desktop Home screen now opens a project menu, and projects opened from the Home screen are registered so they remain available throughout the rest of the app.

— Names the UI action but no further mechanism detail. v1.18.16

02 **Forward-compatible config parsing** IMPROVED 33

Unknown top-level config fields are now silently ignored instead of failing config parsing, making config files forward-compatible.

— Explains behaviour change but no config key names given. v1.18.16

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ AI Model & Data Infrastructure

emisar

1 RELEASE · 2026-08-10 NOTES ↗

An MCP that lets AI tools securely connect to your infrastructure, write IaaS code, debug issues, and assist during incidents - without risking production stability. Built for security teams to approve and infrastructure teams to experience like magic.

> emisar v0.39.0 overhauls runbook versioning with published-release tracking, a diff-and-consent publish flow, and slug-keyed MCP endpoints, while also widening secret redaction and pfSense integration.

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

01 **Slug-keyed runbook MCP endpoints** NEW 78

Adds `get_runbook` and `update_runbook_draft` MCP endpoints keyed by slug, so agents replace only the single unpublished change under the hash they read; publication remains human-only.

— Names endpoints and hash mechanism but not full request shape. v0.39.0

02 **Expanded secret redaction in Runner v0.19.0** IMPROVED 75

Runner `v0.19.0` expands automatic secret redaction to cover connection strings, database URLs, key-derivation inputs (`salt`, `pepper`), cookie and session signing keys, and passphrase spellings — acting as a safety net beneath per-action redaction declarations.

— Enumerates every redacted category and versioned component. v0.39.0

03 **Publish-time diff and consent requirement** NEW 70

Publishing now shows a diff of changed lines against the exact canonical text whose hash defines the release, requiring explicit consent plus the hash of what you read before an unpublished change can run.

— Mechanism well explained though UI path not given. v0.39.0

04 **`not_live` error for stale release runs** IMPROVED 65

Attempting to run a named older release now returns `not_live` instead of silently dispatching current content.

— Names exact error code returned. v0.39.0

05 **pfSense resolver, NTP, WireGuard reads and DHCP action** NEW 65

pfSense gains resolver, NTP, and WireGuard peer reads (secrets excluded), plus a DHCP reservation action staged for operator approval.

— Names the exact read surfaces and the new staged action. v0.39.0

06 **Runbook versioning tracks published releases** IMPROVED 60

Runbook versioning now tracks published releases (e.g. Run v3) rather than counting every save, with a single in-flight draft per runbook instead of unbounded version rows.

— Clear before/after but no direct command to invoke it. v0.39.0

07 **Behavior harness real-system test coverage** IMPROVED 60

Behavior harness now runs tests against a real service manager (PID 1), a per-case Docker daemon, iptables in its own namespace, and a real dpkg database for install, remove, and autoremove; uncovered cases are recorded with explicit reasons.

— Internal test infrastructure, not directly user-actionable. v0.39.0

08 **Automatic runbook version migration on upgrade** BREAKING 60

The portal migration collapses all runbook version rows into a single runbook per slug, renumbers published versions into releases, and repoints execution history — this runs automatically before the instance serves on upgrade.

— Describes migration scope but no manual rollback or opt-out steps. v0.39.0

THINNER COVERAGE BELOW

09 **Execution snapshots for audit integrity** NEW 50

Executions now snapshot the definition they dispatched at start time, so later edits, publishes, or deletes cannot rewrite audit history.

— Explains behavior but no interface to act on directly. v0.39.0

└──▷ BREAKING ON UPGRADE

! The portal migration collapses all runbook version rows into a single runbook per slug, renumbers published versions into releases, and repoints execution history — this runs automatically before the instance serves on upgrade.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Daytona ⓘ 1 RELEASE · seen 2026-08-10 NOTES ↗

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

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

01 **Sandbox TTL (time-to-live) support** NEW 35

Adds sandbox TTL support, enabling automatic expiry control over sandbox lifetimes.

— Names the mechanism but no config key or command shown. `snapshot-20260810`

└──▷ **BREAKING ON UPGRADE**

! Minimum supported Python version raised to 3.10; setups running Python below 3.10 will break on upgrade.

WAS THIS USEFUL? 👍 👎

ai-safe2-framework

2 RELEASES · 2026-07-13 → 2026-07-19 NOTES ↗

The Universal Governance, Risk, Compliance (GRC) Operating System with Integrated Security for Agentic AI, Non-Human Identities, and Swarm Governance. AI SAFE² + AI Sovereignty Maturity Model (AISM) [Dual License: MIT + CC-BY-SA]

> AI SAFE2 v3.0 shipped a Lovable-specific sovereign runtime that closes a subagent file-access exfiltration gap and tightens credential-leak detection, alongside an entirely new CP.11 Unbiased AI Standard defining 27 controls, an 8-class bias taxonomy, published bias-test thresholds, and a GSAR 552.239-7001 compliance evidence package.

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

01 17 named AI SAFE2 v3.0 controls, 21/21 smoke-tested NEW

85

Enforces 17 named AI SAFE2 v3.0 controls — `P1.T1.2`, `P1.T1.10`, `P1.T1.4_ADV`, `P1.T1.1`, `P1.T2.5`, `P1.T2.6`, `S1.3`, `S1.5`, `S1.6`, `P2.T3.1`, `A2.5`, `P3.T5.5`, `P4.T7.1`, `M4.5`, `E5.1`, `CP.4`, `CP.10` — verified at 21/21 smoke-test passing.

Run the full sovereign baseline smoke test to verify all 21 enforcement controls pass before deploying the runtime in your environment.

```
$ cd examples/lovable-sovereign-runtime &&
PYTHONPATH=enforcement python3 smoke_test.py
```

— Full control list plus a runnable smoke-test command

2026-07-18_lovable-sovereign-runtime

02 **CP.11 Unbiased AI Standard: normative spec and 27-control system**

NEW 80

Adds `STANDARD.md` (Sections 1–8) as the normative specification for CP.11, covering scope, definitions, requirements, testing, evidence, and independent review criteria, backed by `controls/uas-controls-system.md` defining 27 controls: 14 system-side controls across pillars P1–P5 and NEXUS, 8 human-side controls sourced from the Cognitive Sovereignty Framework Domains 1–6 (`UAS-H1` through `UAS-H8`), and 5 cross-domain bridge controls. This is the first cross-pillar control to formally integrate the Cognitive Sovereignty Framework (CSF), with CTSS scoring feeding directly into the CP.11 evidence package.

— Names spec files and control breakdown, no runnable example given

`2026-07-13_Unbiased_AI_Standard_(UAS)`

03 **Bias test protocol with published pass thresholds** NEW

75

Adds `testing/bias-test-protocol.md` specifying 8 test types with published pass thresholds: an 85% neutrality classification floor, Cohen's kappa ≥ 0.7 inter-rater minimum, and a p > 0.05 demographic parity requirement across matched framings.

— Exact numeric thresholds named, no runnable test command shown

`2026-07-13_Unbiased_AI_Standard_(UAS)`

04 **Subagent file-access gating closes LV-SUBAGENT gap** NEW

70

Adds a scan_subagent_file_access() enforcement gate that intercepts file path lists before Lovable subagents read them, blocking access to credential files such as `.env.production` by path pattern to close the subagent exfiltration gap. Ships as the `examples/lovable-sovereign-runtime` enforcement package covering a new sixth attack surface, LV-SUBAGENT, that was missed by the prior five-surface design.

— Names function and path pattern; no direct runnable example attached

`2026-07-18_lovable-sovereign-runtime`

05 **8-class bias taxonomy including foreign regulatory embedding**

NEW 70

Adds `taxonomy/bias-taxonomy.md` defining an 8-class bias taxonomy, `UAS-B1` through `UAS-B8` , each mapped to a pillar source, CSF domain, and GSAR trigger — including `UAS-B8` (Foreign Regulatory Embedding), the first published methodology for detecting model modifications made to satisfy foreign regulators.

— Named file and class range, but no usage example to run

`2026-07-13_Unbiased_AI_Standard_(UAS)`

06 **GSAR 552.239-7001 compliance evidence package** NEW 70

Adds `compliance/gsar-552-239-7001-mapping.md` and a five-document due-diligence evidence package mapped paragraph-by-paragraph to GSAR 552.239-7001 sections (j)(1) and (f)(7), including an attestation document at `compliance/due-diligence-attestation.md`.

— Exact file paths and regulation sections named, no example given

2026-07-13_Unbiased_AI_Standard_(UAS)

07 **Workspace-wide trust boundary rules via Lovable Knowledge** NEW 63

Adds `workspace-knowledge/ai-safe2-workspace-knowledge.md`, a no-code drop-in pasted into Lovable → Settings → Knowledge → Workspace knowledge, that applies AI SAFE2 trust boundary rules to every future agent session workspace-wide.

Apply workspace-wide AI SAFE2 trust boundary rules to every Lovable agent session — no code integration required.

> 📍 In the Lovable console, go to Settings › Knowledge › Workspace knowledge and paste the contents of `workspace-knowledge/ai-safe2-workspace-knowledge.md`.

— Exact UI path and file given, no code integration needed

2026-07-18_lovable-sovereign-runtime

THINNER COVERAGE BELOW

08 **Expanded credential and exfiltration detection patterns** IMPROVED 56

Expands detection coverage to block the `sk-live-*` API key format (hyphen in key body), natural-language destructive plan steps such as "delete all existing user records", and Response.json(process.env) (Web API environment variable exfiltration).

— Names exact patterns blocked but gives no way to invoke or verify

2026-07-18_lovable-sovereign-runtime

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 10, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe