THE CYBER TOOLCHAIN

NO. 056

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 25, 2026 · EVERY DAY

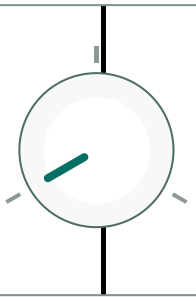
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 31 tools

OFFENSIVE SECURITY

vim ? help

◆ Reconnaissance & OSINT

Maigret

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗ CODE ↗

Collect a dossier on a person by username from 3000+ sites

> Maigret v0.6.5 expands its username-search database with dozens of new sites and fixes two long-standing sites-database problems: a credential leak and a permission crash on read-only installs.

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

01 **Fixed PermissionError on read-only installs** IMPROVED

73

Removed the unconditional post-run write of the sites database to the package directory, so Maigret no longer exits with `PermissionError` on read-only installations such as distro packages, snaps, the Nix store, or containers; the only remaining write now goes to `~/.maigret`.

– Names exact error, affected environments, and the new write path v0.6.5

THINNER COVERAGE BELOW

02 **Expanded site coverage across batches** NEW

50

Adds 78+ new sites to the username-search database across multiple batch additions, including `VNOJ` (Vietnamese competitive programming), `Juejin`, `NetEase Music`, `Coddy`, and `Codédex`, plus 6 additional Chinese sites with presence markers to expand coverage for OSINT investigations targeting Chinese-language platforms.

– Names specific sites and counts but no usage mechanism given v0.6.5

03 **Runtime tokens excluded from sites database** IMPROVED

38

Runtime activation tokens are now kept out of the sites database, preventing credential leakage in the distributed `data.json`.

– States the fix and file name but no further mechanism v0.6.5

WAS THIS USEFUL? 👍 👎

Nettacker

Feedback

1 RELEASE · 2026-08-24

NOTES ↗

Automated Penetration Testing Framework - Open-Source Vulnerability Scanner - Vulnerability Management

> Nettacker 0.4.1 adds SARIF and DefectDojo-compatible JSON output, expands scan configuration and Web UI controls, ships several new scanner modules, and moves its runtime data directory with automatic migration.

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

01 **Runtime data directory moved to `.nettacker/data`** BREAKING 70

Runtime data, including the SQLite database and scan results, has moved from `.data` to `.nettacker/data`, with automatic database migration on upgrade. Any scripts or tooling referencing the old `.data` path must be updated.

— Names exact old and new paths and migration behaviour 0.4.1

THINNER COVERAGE BELOW

02 **New scanner modules and coverage expansion** NEW 55

Adds new scanners for CUPS, configuration files, SMB brute-force, Jenkins, SmarterMail, CrushFTP patch-date, and Adobe AEM patch-date, alongside expanded WordPress plugin coverage and refreshed bundled discovery wordlists.

— Names seven scanners plus wordlist refresh, no run commands 0.4.1

03 **Web UI scan and results improvements** IMPROVED 53

The Web UI now exposes service-discovery, port-exclusion, and custom-header options for scans, adds browser-based file uploads for targets and wordlists, and adds searchable profiles, scan methods, and results.

— Groups three UI additions with named surfaces, no exact paths 0.4.1

04 **SARIF and DefectDojo JSON output formats** NEW 50

Scan results can now be exported as SARIF and as DefectDojo-compatible JSON, in addition to existing formats.

— Names two formats but no export command or flag 0.4.1

05 **Expanded scan configuration options** NEW 43

Scan configuration gains port exclusions, HTTP scheme selection, custom request headers, and sensitive-header filtering.

— Lists four options but no flag names or syntax 0.4.1

06 **Custom wordlists for brute-force and discovery** NEW

40

Adds custom wordlists and directory-enumeration base paths for brute-force and discovery modules.

— Describes capability without config syntax 0.4.1

07 **crt.name subdomain discovery source** NEW

35

Adds `crt.name` as a source for subdomain discovery.

— Names the source but no usage detail 0.4.1

08 **FreeBSD and ARM platform support** NEW

35

Adds FreeBSD support and ARM container images.

— States platforms added, no further mechanism 0.4.1

└──▷ **BREAKING ON UPGRADE**

! Runtime data (SQLite database and results) has moved from `.data` to `.nettacker/data`; any scripts or tooling that reference the old `.data` path must be updated.

WAS THIS USEFUL? 👍 👎

bbot

Feedback ⓘ 1 RELEASE · 2026-08-24 NOTES ↗

The recursive internet scanner for hackers.

> BBOT 3.0.2 ships a new js-audit preset for hunting leaked credentials and vulnerable JavaScript libraries, alongside HTTP decoding hardening and increased worker throughput.

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

01 js-audit preset for client-side JavaScript auditing NEW

88

A new `js-audit` preset (shipped at `bbot/presets/web/js-audit.yml`, and documented as living at `~/.bbot/presets/web/js-audit.yml`) chains the `trufflehog`, `retirejs`, `badsecrets`, `http`, `robots`, and `wayback` modules to detect leaked credentials and vulnerable JavaScript libraries in client-side code. It enables the `only_verified` and `include_sitemap` options, runs trufflehog against HTTP bodies streamed via stdin while including the source URL on findings, and sets `_module_threads` to 2.

Scan a target web app for leaked credentials and outdated JavaScript libraries using the new built-in preset.

```
$ bbot -t example.com -p js-audit
```

Audit JavaScript files for secrets (e.g. API keys, tokens) across a target's web surface using the new `js-audit` preset.

```
$ bbot -t www.evilcorp.com -p js-audit
```

— Names exact modules, config path, options and a runnable command v3.0.2

THINNER COVERAGE BELOW

02 decode-error tag on undecodable HTTP responses IMPROVED

53

`HTTP_RESPONSE` events now get a `decode-error` tag when a response body cannot be decoded, preventing undecoded content from being treated as valid body data.

— Explains mechanism and effect but no reader action v3.0.2

03 Increased worker recycling threshold IMPROVED

50

Process pool `max_tasks_per_child` was raised from 25 to 250, enabling more sustained parallel workloads before worker recycling.

— Gives exact before/after numbers but no config exposed v3.0.2

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Bugcrowd is a crowdsourced security platform connecting organizations with ethical hackers to identify vulnerabilities through bug bounty programs.

> Bugcrowd updated its Request a Response (RaR) workflow so submissions no longer auto-expire, giving researchers the ability to request expiration removal and giving customers new tracking and filtering tools for these requests.

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

01 Request a Response expiration removal and tracking NEW

67

Researchers can now use Request a Response Expiration Removal to ask that a response expiration be removed from a submission, and RaRs no longer auto-expire. Each RaR now displays the date it was raised so teams can track how long it has been open, and the inbox adds filters for Active requests to Customers (outstanding requests) and Resolved requests to Customers (handled ones).

Quickly surface all outstanding researcher requests that need your team's attention using the new inbox filter.

📍 In the Bugcrowd inbox, open the filter panel and select Active requests to Customers to see all open RaRs awaiting a response.

— Names exact filter labels and UI steps but lacks deeper mechanism detail.

Request a Response Expiration Removal (…

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ 1 RELEASE · 2026-08-20 NOTES ↗

Documentation for the Dreadnode AI security platform — self-hosting guides, platform concepts, SDK reference, and CLI documentation.

> Dreadnode's latest release adds Kubernetes Gateway API routing for on-prem installs, a full Caido TypeScript SDK skill for web-security agents, three new attack techniques, and several usability improvements to session management and Worlds data export.

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

01 Caido TypeScript SDK skill for web-security agents NEW

80

A new `caido-mode` skill in the web-security capability exposes the full Caido TypeScript SDK from within an agent, including Match & Replace rules, replay sessions, HTTPQL history search, findings management, and curl-through-Caido testing.

— Names the skill and every SDK capability it exposes.

Kubernetes Gateway API support for on-p…

02 webhook.site API key for callback tool rate limits IMPROVED

75

The web-security callback tool now supports a `WEBHOOK_SITE_API_KEY` environment variable, letting a paid webhook.site account bypass rate limits while preserving anonymous and interactsh fallback.

Use a paid webhook.site account in the callback tool to avoid anonymous rate limits during web-security agent runs.

```
$ export WEBHOOK_SITE_API_KEY=your-api-key-here
```

— Exact env var and runnable export command are given.

Kubernetes Gateway API support for on-p…

03 Kubernetes Gateway API routing for on-prem installs NEW

65

On-prem deployments can now use Kubernetes Gateway API (HTTPRoute) routing instead of relying on Ingress, enabling installation on Istio, Envoy Gateway, and Cilium stacks.

— Names the resource and supported gateway stacks but no full config steps.

Kubernetes Gateway API support for on-p…

04 **ATIF dataset and replay export in Worlds UI** NEW 65

The Worlds UI now supports downloading the full ATIF dataset from Worlds → Jobs, and exporting the normalized replay payload as JSON from the Trajectories page.

— Exact UI navigation and export format given. Kubernetes Gateway API support for on-p…

THINNER COVERAGE BELOW

05 **New web-security attack skills** NEW 55

Three new attack skills were added to the web-security capability: SQL-injection-to-RCE escalation, MathJax/KaTeX macro injection, and an ORB/XS-Leak browser side-channel attack via service-worker fetch proxy.

— Names three distinct attacks but no invocation details given.

Kubernetes Gateway API support for on-p…

06 **Typed delete confirmation for sessions** NEW 50

Bulk and single-row session deletes now require typing a `DELETE-{count}` challenge phrase to confirm, guarding against accidental permanent deletion.

— Names exact confirmation format but limited further detail.

Kubernetes Gateway API support for on-p…

07 **Bulk freeze for agent sessions** NEW 45

Multiple agent sessions can now be frozen at once from the Agent Sessions tab, with a confirmation dialog shown before sessions are affected.

— Clear UI location but no further mechanism detail.

Kubernetes Gateway API support for on-p…

08 **Media previews in output panel** IMPROVED 25

Image, audio, and video previews are now embedded directly in the output panel alongside existing field data.

— Thin description with no further mechanism or scope.

Kubernetes Gateway API support for on-p…

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Aikido Security

Feedback ⓘ 5 RELEASES · SEEN 2026-08-25

NOTES ↗

Aikido Security is a developer-focused security platform that scans code, infrastructure, and dependencies for vulnerabilities and misconfigurations.

> Aikido shipped a wave of agentic security capabilities this quarter — pentest-grade source code reasoning with Code Audit, code-based DSPM via Data Exposure Audit, and LLM-driven triage agents for secrets and CVEs — alongside GA GCP VM scanning, an expanded cross-platform Device Protection agent (now including Linux), and a set of new bulk-configuration APIs and reporting formats.

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

01 **Device Protection agent for endpoints, now including Linux**

NEW 85

A new Device Protection agent provides real-time malware blocking across npm, PyPI, Maven, and NuGet package managers (powered by Aikido Intel), extension governance covering VS Code, OpenVSX, Chrome, and skills.sh, a request-and-approval workflow where blocked installs trigger a one-click exception request reviewable in the Aikido admin Inbox, group-based policies for per-team rules, and continuous device inventory. The agent runs silently and only surfaces desktop notifications when a package is blocked or held for approval. Device Protection for Linux (beta) extends this to Linux desktops and servers, distributed as `.deb` and `.rpm` packages.

.
.
.
.

— Names package managers, extension surfaces, and Linux package formats

#54 - Device Protection for Linux and V… | 🚀 Protect Company Devices From Supply …

02 **Data Exposure Audit (code-based DSPM)** NEW 81

A new Data Exposure Audit feature performs code-based DSPM using LLM agents to trace sensitive fields from sources (request bodies, DB reads, config) to sinks (logs, third-party APIs, LLM prompts, caches, queues, storage) without connecting to production data stores. It detects credentials in URLs, API over-exposure (endpoints returning more fields than the caller is entitled to, missing tenant ownership checks), unsafe logging/exports of unredacted PII and secrets, over-broad third-party and AI sharing of customer data, weak storage (unencrypted or only base64/MD5/SHA-1 protected), incomplete deletion flows, and over-permissioned IAM roles/OAuth scopes

and unauthenticated data stores. It runs against already-connected repositories with no new data source to onboard, and each audit consumes credits from the Aikido Wallet.

— Names every detection category and mechanism; no exact command given

🚀 DSPM: Find data exposure risks witho…

03 **Code Audit: agentic source-code security reasoning** NEW

72

Code Audit runs agentic, pentest-grade security reasoning directly on source code to find logic flaws (e.g. cross-tenant data leakage) that pattern-based scanners miss, requiring no staging URL, traffic replay, or agent deployment. AutoFix generates remediation PRs from Code Audit findings with a few clicks alongside root cause analysis and remediation guidance, and it supports hard-to-test-by-URL codebases including mobile apps, smart contracts, desktop apps, embedded software, and libraries. It is priced in credits shown upfront before the audit starts, based on codebase size and complexity.

— Explains mechanism and pricing but no exact command shown

🚀 Code Audit: Pentest-grade reasoning …

04 **CVE Exploitability Analysis agent, now for containers**

IMPROVED

65

A new CVE Exploitability Analysis agent uses LLM reasoning over how code uses an affected package to evaluate whether a CVE is actually reachable and exploitable, and can act autonomously — upgrading, downgrading, snoozing, or ignoring findings — based on rules teams control, with each run logged with a short rationale for review or override. CVE Exploitability assessment now also supports containers.

— Explains behaviour and autonomy but no config surface named

#54 - Device Protection for Linux and V…

🚀 Extreme Noise Reduction for Secrets …

05 **Secrets AutoTriage agent** NEW

65

A new Secrets AutoTriage agent runs after entropy detection to assess whether a discovered secret is a real, exposed credential — evaluating liveness, scope, and surrounding code context — then applies defined rules or LLM reasoning for ambiguous cases. It can act autonomously (upgrading, downgrading, snoozing, or ignoring findings) based on rules teams control, and each run is logged with a short rationale for review or override.

— Describes triage logic and autonomy without a concrete config surface

🚀 Extreme Noise Reduction for Secrets …

THINNER COVERAGE BELOW

06 **GCP VM Scanning reaches general availability** IMPROVED

50

GCP VM Scanning is now generally available with agentless setup requiring no SSH keys or agents, completing agentless VM scanning coverage across AWS, Azure, and GCP.

— States scope and mechanism but no setup steps given

#54 - Device Protection for Linux and V…

07 **OSCAL FedRAMP DAST report format** NEW

50

Adds an OSCAL FedRAMP DAST report format under Reports.

— Names format and location but no generation steps

#54 - Device Protection for Linux and V…

08 **Ignore controls for issues and license findings** NEW

45

Adds the ability to request an ignore of an issue on a PR/MR with admin approval, and the ability to ignore license issues by path.

— Two thin ignore workflows named without further mechanism

#54 - Device Protection for Linux and V…

09 **Zen Firewall agent expands language support** NEW

35

Zen Firewall agent releases are now available for Dotnet, Java, PHP, Ruby, Golang, and Node.

— Names languages but no setup or usage detail #54 - Device Protection for Linux and V…

10 **Bulk GitHub PR configuration API** NEW

35

A new API allows bulk updating of GitHub PR configuration across repositories.

— Named API surface but no endpoint path or method given

#54 - Device Protection for Linux and V…

11 **Package Health adds NuGet and JetBrains support** NEW

30

Package Health now supports NuGet and JetBrains packages.

— One-line addition with no mechanism described #54 - Device Protection for Linux and V…

12 **Hardened container image request API** NEW

30

A new API allows requesting a hardened container image.

— Named API surface but no endpoint path or parameters given

#54 - Device Protection for Linux and V…

13 **New VSCode IDE extension release** IMPROVED 15

A new release of the VSCode IDE extension was published.

— Bare mention of a release with no changelog detail

#54 - Device Protection for Linux and V…

WAS THIS USEFUL?  

JFrog

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

JFrog is a DevOps platform providing software supply chain management, artifact management, and continuous integration/delivery automation tools.

> JFrog's biggest move this window is Agent Package Resolution (preview), which routes AI coding agent package downloads through Artifactory governance automatically, alongside a smaller update giving the Catalog service bulk governance automation via GraphQL.

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

01 Agent Package Resolution for AI coding agents NEW

91

A preview feature of the JFrog Agent Plugin routes AI coding agent dependency downloads (Claude Code, Cursor) through JFrog Artifactory instead of public registries, covering npm, PyPI, Maven, Go, Docker Hub, Helm, and NuGet. Session steering automatically supplies the Artifactory URL for each governed package type at the start of an agent session with no per-project configuration, while the new `jf setup` command persistently binds a local package manager to an Artifactory repository across sessions and tools. JFrog Curation policy enforcement is extended server-side to all agent-driven downloads — blocking malicious packages, unapproved licenses, and policy-violating packages regardless of whether session steering or `jf setup` is active — and covers indirect/transitive installs such as postinstall scripts, `pip install -r requirements.txt`, and `npm ci`, producing complete audit trails for every agent-driven download, block, and policy decision alongside existing human-driven build logs.

One-time project setup: bind your local package managers to Artifactory so that every subsequent agent session — and any indirect installs — resolves through the governed repository without manual configuration.

```
$ jf setup
```

Enable Agent Package Resolution for your team via the JFrog Agent Plugin on GitHub so Claude Code or Cursor sessions automatically receive the correct Artifactory URLs at startup.

> 📍 In the JFrog Platform, install the JFrog Agent Plugin from GitHub, then under Agent Package Resolution settings choose which package types (npm, PyPI, Maven, Go, Docker, Helm, NuGet) are governed, and enable zero-touch setup to push Artifactory URLs to agent sessions automatically.

— Names exact command, package types, and policy mechanism

Coding Agents Just Reopened Your Softwa…

THINNER COVERAGE BELOW

02 Bulk package governance via Catalog GraphQL API NEW

38

JFrog Catalog now supports bulk automation via its GraphQL API, enabling package governance at scale.

— Names the API but no endpoint, schema, or usage detail product docs

WAS THIS USEFUL? 👍 👎

GitLab

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

The intelligent orchestration platform for DevSecOps, enabling teams and agents to ship trusted software at enterprise scale.

> GitLab Dedicated now offers fully managed, single-tenant Hosted Runners, removing the need for customers to provision and maintain their own CI runner fleet.

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

01 **Hosted Runners for GitLab Dedicated** NEW 89

GitLab Dedicated gains a fully managed, single-tenant CI runner fleet isolated per customer, eliminating the need to provision or maintain your own runners. It supports machine sizes `linux-small-amd64` (default, 2 vCPU / 8 GB), `linux-medium-amd64` (4 vCPU / 16 GB), `linux-large-amd64` (8 vCPU / 32 GB), `linux-xlarge-amd64` (16 vCPU / 64 GB), and `linux-2xlarge-amd64` (32 vCPU / 128 GB) on Linux x86-64 and Linux Arm64. Each CI job runs in a newly provisioned, short-lived VM deleted after completion for job-level process isolation; secure connectivity to internal artifact registries, secret managers, and deploy targets is available via AWS PrivateLink. Pricing is consumption-based via GitLab Credits with usage visible in the GitLab Credits Dashboard, the fleet is self-serviceable through Switchboard (the GitLab Dedicated management console), and the offering carries a 99.9% uptime SLA with autoscaling to handle bursty CI workloads.

Target a specific hosted runner size in a GitLab Dedicated pipeline job instead of self-managed runner tags.

```yaml
build:
  tags:
    - linux-large-amd64
  script:
    - make build
```

Use the default small hosted runner when no size preference is specified — the `linux-small-amd64` tag is applied automatically.

```yaml
test:
  tags:
    - linux-small-amd64
  script:
    - make test
```

— Names sizes, SLA, billing and console; includes runnable tag examples.

Scale software delivery pipelines in is…

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC." name="description"/> <meta content=

> Endor Labs extended its AI SAST engine with full C language support, combining buildless scanning with deterministic program analysis and LLM reasoning to catch multi-function memory bugs.

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

01 **C language support in AI SAST** NEW 58

Endor Labs AI SAST now covers C with buildless scans that run in the IDE and on PRs without requiring a compiler hook or reproducible build. It pairs deterministic call-graph and reachability analysis with LLM reasoning to trace data flow across functions and files, catching buffer overflows, use-after-free, and integer overflows that span multiple functions; each finding includes a call path, a working exploit, and a suggested fix, matching the evidence format used for other languages. C SAST runs alongside C SCA in the same platform to cover first-party code and its open-source dependencies together, and the scanning runs as part of AURI by Endor Labs, an independent security harness outside the coding agent that enforces policy across any agent, model, or CI stage.

— Describes mechanism and scope well but no concrete flags, commands, or limits given

Why Endor Labs AI SAST for C finds what…

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

A fully-managed, enterprise-scale solution to control, secure, and distribute software packages and containers.

> Cloudsmith's biggest addition is Early Access Nix binary cache repository support with Ed25519 signing and NixOS upstream channel proxying, alongside upstream proxying for hardened container images and curated libraries, and a new Broadcasts feature that publishes repositories to a branded public URL with aggregated usage analytics.

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

01 Nix binary cache repository support (Early Access) NEW

95

Cloudsmith now hosts NAR and narinfo files over the standard Nix HTTP binary cache protocol at `https://nix.cloudsmith.io/OWNER/REPOSITORY`. Packages can be published with `nix copy --to 'https://nix.cloudsmith.io/OWNER/REPOSITORY' --no-check-sigs` and fetched with `nix copy --from 'https://nix.cloudsmith.io/OWNER/REPOSITORY' STORE_PATH` once the repository is added as a substituter via `extra-substituters` and `extra-trusted-public-keys` in `nix.conf` (e.g. `/etc/nix/nix.conf` or `~/.config/nix/nix.conf`). Private repositories support entitlement tokens embedded directly in the substituter URL (`https://token:TOKEN@nix.cloudsmith.io/OWNER/REPOSITORY`) or via a `netrc` file referenced with `netrc-file`. Each repository gets a per-repository ED25519 signing key, with Cloudsmith re-signing every narinfo on every read to allow immediate key rotation and multiple simultaneously active keys. Upstream proxying and caching against NixOS release channels (e.g. `https://channels.nixos.org/nixos-26.05`) is also supported, with each upstream mapped to exactly one channel via the Proxy URL field; this is currently web-app only, with API-based upstream management not yet available.

Publish a locally built Nix package and its full dependency closure to a private Cloudsmith Nix repository.

```
$ STORE_PATH="$(readlink -f result)"
nix copy --to 'https://nix.cloudsmith.io/OWNER/REPOSITORY' --no-check-sigs "$STORE_PATH"
```

Configure a Cloudsmith Nix repository as a system-wide substituter so `nix build` transparently fetches cached packages instead of building locally.

```ini
extra-substituters = https://nix.cloudsmith.io/OWNER/REPOSITORY
extra-trusted-public-keys = KEY_NAME:PUBLIC_KEY
```

Configure your Nix client to use a Cloudsmith repository as a substituter for a proxied NixOS release channel, with `cache.nixos.org` as a fallback for closure members Cloudsmith cannot index.

```ini
extra-substituters = https://nix.cloudsmith.io/OWNER/REPOSITORY
https://cache.nixos.org
extra-trusted-public-keys = KEY_NAME:PUBLIC_KEY cache.nixos.org-
1:6NCHdD59X431o0gWypbMrAURkbJ16ZPMQFGspcDShjY=
```

Create a Nix upstream in Cloudsmith pointing at a specific NixOS release channel; create a separate upstream for each channel you want to proxy.

> 📍 In the Cloudsmith web app, go to your repository › Upstreams › Add Upstream › Nix, set the URL to `https://channels.nixos.org/nixos-26.05`, configure authentication (or select `None`), then save. Repeat with `https://channels.nixos.org/nixos-25.11` for a second channel.

— Full mechanism, config keys and runnable commands given; API management still missing.

New format: Nix, with upstream support …

---

02 **Broadcasts: public branded repository URLs with analytics**

NEW

70

Broadcasts exposes a standard Cloudsmith repository to the public through a customizable, branded URL instead of a separate public repository, supporting both a default broadcast URL and workspace custom domains for browser browsing and native tooling pulls (Docker, npm, Maven, IDE, CI/CD). A dashboard Overview section shows aggregated analytics across broadcasting repositories — total requests per billing cycle, usage per package, recently downloaded packages, recently created broadcasts, a breakdown of requests by public/private/open source packages, and entitlement tokens consumed in the current billing cycle — plus a per-package toggle between storage `Usage` and `Requests` views. Broadcasting can be enabled at repository creation via a toggle, or toggled on existing repositories via `Start broadcasting` / `Stop broadcasting` in the repository page dropdown.

Enable broadcasting on an existing repository so end-users can pull packages using native tooling from your workspace's public broadcast URL.

> 📍 In the Cloudsmith dashboard, go to the main page for your repository, expand the dropdown next to `Not broadcasting`, and click `Start broadcasting`.

— Clear UI path given, but no API or config key for automation. product docs

THINNER COVERAGE BELOW

03 **Upstream proxying for hardened images and curated libraries**

NEW 40

Adds upstream proxying for library registries (Chainguard Libraries, RapidFort Curated Libraries), letting hardened or filtered open-source package replacements be proxied through a Cloudsmith repository, and for hardened image registries, providing minimal, vulnerability-reduced container images from vendor-maintained sources.

:

— Named integrations only; no config, endpoint or setup steps given. product docs

WAS THIS USEFUL? 👍 👎

Mend.io Mend SAST

Feedback ⓘ 1 RELEASE · 2026-08-24 NOTES ↗

Traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend SAST 26.8.1 adds a probability risk factor per data flow for more granular triage, expands CWE detection coverage for JavaScript/TypeScript and Rust, and introduces an opt-in option to share project data so AI can generate customized detection rule candidates.

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

01 **New CWE detections for JS/TS and Rust** NEW 60

Adds detection for `CWE-99` (Improper Control of Resource Identifiers / Client Resource Injection) in JavaScript and TypeScript, plus detection for `CWE-319` (Cleartext Transmission of Sensitive Information) and `CWE-321` (Use of Hard-coded Cryptographic Key) in Rust.

— Named CWE IDs and languages but nothing for reader to configure

Version 26.8.1 (24-August-2026)

THINNER COVERAGE BELOW

02 **Probability risk factor per data flow** NEW 50

Mend SAST now displays a probability risk factor next to each data flow in finding details within the Mend Platform, enabling more granular triage.

— Names the UI location but no config or API detail Version 26.8.1 (24-August-2026)

03 **Opt-in AI rule customization** NEW 40

Adds an opt-in option to share additional project data with Mend.io so AI can generate candidates for customizing and tailoring detection rules to your organization.

— Describes the opt-in mechanism but no named toggle or setting

Version 26.8.1 (24-August-2026)

WAS THIS USEFUL? 👍 👎

Veracode

Feedback

1 RELEASE · 2026-08-24

BLOG ↗

Veracode is a cloud-based application security platform that scans code for vulnerabilities and provides remediation guidance.

> Veracode's biggest update this window is GitHub Workflow Integration v2.0.0/v2.1.1, which lets teams install, configure, schedule, and uninstall the GitHub integration directly from the Veracode Platform, alongside expanded Java API Wrapper endpoint coverage, new `veracode.yml` configuration options across Azure DevOps/GitHub/GitLab integrations, and two upcoming deprecations (the Veracode Flaw Importer task and Jira Server/Data Center 9.x support).

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

01 Java API Wrapper endpoint coverage and reliability

IMPROVED

75

Java API Wrapper v26.3.19.0 adds page size support to the Platform Applications API (`/appsec/v1/applications`) and introduces support for the paginated Sandbox REST API. Java API Wrapper v26.2.18.0 adds support for the Veracode Platform Findings REST API to access open and mitigated application findings, and implements a configurable retry mechanism for the Platform Applications API (`appsec/v1/applications`) to handle transient failures and timeouts.

— Endpoints and wrapper versions named but no code sample given

Integration updates (2026-08-24)

02 GitHub integration managed from Veracode Platform

NEW

70

GitHub Workflow Integration v2.0.0 enables installing, configuring, and updating credentials for the Veracode GitHub Integration directly from the Veracode Platform, eliminating manual GitHub Marketplace steps, and adds the ability to schedule Static Analysis, SCA, and Infrastructure as Code (IaC) scans for GitHub organizations from the Platform. GitHub Workflow Integration v2.1.1 adds a new uninstallation workflow on the Veracode Platform to remove the Veracode app from GitHub organizations.

— Names versions and scan types but only points to a Platform UI flow

Integration updates (2026-08-24)

03 Azure DevOps Workflow Integration scheduling and config options

NEW

70

Azure DevOps Workflow Integration v2.4.3 adds a custom agent pool name parameter in `veracode.yml` to specify the agent pool used for pipeline execution; v2.3.10 adds `analysis_branch` support so scheduled scans run against a specified branch instead of the default; and v2.3.9 adds the ability to schedule scans for every organization where the integration is installed, based on `veracode.yml` configurations.

Run scheduled scans against a specific analysis branch rather than the default branch in Azure DevOps.

```yaml
analysis_branch: main-security-review
```

— Names each config key and file, with a runnable example

Integration updates (2026-08-24)

---

04 **VeracodeFlawImporter task replaces Veracode Flaw Importer**

65

BREAKING

Veracode Azure DevOps extension v3.35.0 replaces the deprecated `Veracode Flaw Importer` task (with spaces) with the new `VeracodeFlawImporter` task for compatibility with a future Microsoft TFX packaging library upgrade; the old task name is scheduled for removal in a future release.

— Exact old and new task names given, migration is clear Integration updates (2026-08-24)

---

05 **Debug option added to ADO, GitHub, and GitLab integrations**

NEW

60

Adds the `debug` configuration option in `veracode.yml` for the Azure DevOps Workflow Integration v2.3.6, GitHub Workflow Integration v1.0.6, and GitLab Workflow Integration v2.1.4 to enable debugging and troubleshooting.

Enable debug logging in a Veracode workflow integration to troubleshoot scan pipeline issues.

```yaml
debug: true
```

— Exact config key and file given with a runnable example

Integration updates (2026-08-24)

---

06 **Jira Server/Data Center 9.x support ending** DEPRECATED

60

Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026; upgrade to Jira 10 or later to continue using a supported version.

— Exact end date and required version given for migration

Integration updates (2026-08-24)

THINNER COVERAGE BELOW

07 **GitHub runner image updated to scm-packaging-5.0.0** IMPROVED

55

GitHub Workflow Integration v2.1.1 updates the default runner image to `scm-packaging-5.0.0`, adding Alpine v3.24.0, Ubuntu v26.04, .NET v10, PipEnv v2026.6.2, libicu v78, Node.js v22.22.1, and Python v3.14.4.

— Exact tool versions given but no action beyond the automatic update

Integration updates (2026-08-24)

08 **Scan Scope setting in Visual Studio plugin** NEW 50

Adds a Scan Scope setting in Veracode Scan for Visual Studio v0.8.0 to scan either the full project directory or only the active solution for faster, more focused results.

— Names the setting and its two modes, no exact UI path Integration updates (2026-08-24)

09 **GitHub autopackaging for Apex, VB6, and Perl** NEW 30

GitHub Workflow Integration v2.0.0 adds autopackaging support for Apex, VB6, and Perl.

— Bare statement of new language support, no usage detail

Integration updates (2026-08-24)

10 **GitHub Veracode Workflow App permissions reduced** IMPROVED

30

GitHub Veracode Workflow App permissions are downgraded to remove administration permissions, preventing the app from altering or removing repositories where it is installed.

— Security-relevant change but no scope or migration steps given

Integration updates (2026-08-24)

11 **Usage analytics in Visual Studio, JetBrains, and VS Code plugins**

NEW 30

Adds usage analytics in Veracode Scan for Visual Studio v0.8.0 and Veracode Scan for JetBrains v1.8.0 and VS Code v1.16.0 to provide insights into plugin adoption, feature usage, and scan activity.

— States purpose only, no mechanism or opt-out detail Integration updates (2026-08-24)

12 **Visual Studio 2026 support** NEW 20

Veracode Scan for Visual Studio v0.7.2 adds support for Visual Studio 2026.

— Bare compatibility statement with no further detail Integration updates (2026-08-24)

└──▷ **BREAKING ON UPGRADE**

! The `Veracode Flaw Importer` task (with spaces in the task name) in the Veracode Azure DevOps extension is deprecated and scheduled for removal in a future release; replace it with `VeracodeFlawImporter`.

! Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026; upgrade to Jira 10 or later to continue using a supported version.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Bright Security

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

Bright is a dynamic application & API security testing platform that offers automation, fast scans, minimal false positives, and comprehensive security testing.

> Bright's August 25, 2026 release introduces Workflows for orchestrating Discovery-plus-Scan pipelines, adds an OWASP Top 10:2025 dashboard widget, overhauls entry-point status filtering in both the Scan Configuration API and Scan by Status UI, and sharpens several vulnerability detectors while dropping the X-XSS-Protection header check.

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

01 entryPointFilter replaces entrypointsStatuses in Scan API

95

BREAKING

The Scan Configuration API gains `entryPointFilter`, supporting filtering by both `securityStatus` and `connectivityStatus` (values include `'ok'`, `'unauthorized'`, `'problem'`, `'unreachable'`) for precise scan scoping. The older `entrypointsStatuses` property is scheduled for deprecation on 22 September 2026, so integrations must migrate to `entryPointFilter` before that date.

— Exact field names, values, and a hard migration deadline Release of August 25, 2026

02 Bright Workflows orchestration pipeline NEW

60

Bright Workflows is a new visual orchestration layer that chains a trigger, Discovery, and Scan into a reusable, schedulable pipeline. It automatically identifies new and changed entry points since the last baseline and passes only those to the Scan.

— Mechanism described but no exact config or command shown Release of August 25, 2026

THINNER COVERAGE BELOW

03 Connectivity status filtering for Scan by Status NEW

55

Scan by Status now supports filtering by Connectivity Status, defaulting to entry points with `OK` status and optionally including `Unauthorized`, `Problem`, or `Unreachable` endpoints alongside the existing Entry Point Status filter.

— Named statuses and default behaviour, UI-only path Release of August 25, 2026

04 OWASP Top 10:2025 dashboard widget NEW

50

A new OWASP Top 10:2025 Risk Overview widget was added to both Organization and Project dashboards, aggregating open issues across selected projects mapped to OWASP categories.

– Clear UI location but no configuration detail Release of August 25, 2026

05 **X-XSS-Protection header check removed** BREAKING 45

The X-XSS-Protection misconfiguration check has been removed from Header Security tests; scanners that previously relied on this finding will no longer receive it.

– Named check and clear effect, no migration path given Release of August 25, 2026

06 **Secret Tokens detection accuracy for Bing Maps and LinkedIn** IMPROVED 40

Secret Tokens detection for Bing Maps and LinkedIn keys now uses tighter patterns that exclude bundled content and unrelated marketing data, reducing false positives.

– Names affected key types, no operable surface Release of August 25, 2026

07 **CORS Misconfiguration detection expanded** IMPROVED 35

CORS Misconfiguration detection now covers origin reflection with credentials, described as the highest-severity CORS bypass scenario.

– Names the scenario but no reader action available Release of August 25, 2026

08 **EXIF Metadata detection broadened** IMPROVED 35

EXIF Metadata detection now covers API responses that return site-root image paths without a leading slash.

– Specific path-format fix, no user-facing control Release of August 25, 2026

09 **CSS Injection out-of-band confirmation improved** IMPROVED 35

CSS Injection out-of-band confirmation now maintains browser context during callback verification.

– Mechanism named but no reader-facing action Release of August 25, 2026

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! The `entrypointsStatuses` property in the Scan Configuration API is scheduled for deprecation on 22 September 2026; migrate integrations to `entryPointFilter` (with `securityStatus` and `connectivityStatus` sub-fields) before that date.

! Removes the `X-XSS-Protection` misconfiguration check from Header Security tests; scanners that previously relied on this finding will no longer receive it.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

GitGuardian SaaS is a cloud-based platform that detects and prevents secrets and sensitive data from being exposed in Git repositories.

> GitGuardian SaaS extended its detection engine beyond code repositories to scan AWS S3 buckets, including secrets buried inside nested archives.

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

01 **AWS S3 bucket scanning for secrets** NEW 83

GitGuardian can now scan AWS S3 buckets using the same detection engine and incident workflow as repositories, surfacing secrets in objects like CI logs, Terraform state, database dumps, `.env` files, and support exports. It connects via OIDC role assumption — GitGuardian assumes a read-only IAM role you create and scope in your own AWS account, with no long-lived AWS keys required — and opens and scans inside archives (e.g. `tar.gz`) nested in S3 objects. You select individual buckets to monitor and assign them to teams like any other source, starting with high-risk buckets such as backup and log archives; scans are scoped to the integration's configured region and S3 Standard storage class by default, with cross-region scanning and other storage classes opt-in to control retrieval and transfer costs. Incidents found in S3 carry the same severity ratings, validity status, playbooks, and notifications as code-repository incidents. Historical (on-demand) scans are available now, with recurrent scanning of new and modified objects coming soon.

Set up the AWS S3 integration in GitGuardian by creating the OIDC provider and read-only IAM role, then registering your account.

> 📍 In the GitGuardian console, go to Integrations › AWS S3 and follow the three-step setup: create the OIDC provider in AWS, create the read-only IAM role scoped to the buckets you want monitored, then enter your AWS account ID and role name in the integration settings.

Connect an AWS account so GitGuardian can scan S3 buckets without handing over long-lived keys.

> 📍 In GitGuardian, go to Integrations › AWS S3, then follow the three-step setup: create the OIDC provider in AWS, create a read-only IAM role scoped to the buckets you want monitored, and register your AWS account ID and role name in the integration settings.

Kick off a historical scan on high-risk buckets (backup archives, log stores) to surface secrets that have never been visible to code scanners.

📍 In GitGuardian, go to Integrations › AWS S3 › your connected account, select the individual buckets to scan (e.g. your backup and log-archive buckets), assign them to the relevant team, and click Trigger historical scan.

— Names OIDC role, storage-class scoping, and UI path but no CLI/API.

Scan your AWS S3 buckets for secrets (2…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ **Firewalls & Network Defense**

Fastly

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

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

> Fastly added a new metric for tracking headless bot traffic across its Historical Stats and Real-Time Analytics APIs.

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

01 **Headless bot request metric in Analytics APIs** NEW

55

The Historical Stats API and Real-Time Analytics API now expose a `bot_edge_requests_headless_count` metric to track edge requests originating from headless bots.

— Names exact metric field but no usage example or query detail.

Add headless bot metric to Historical S…

WAS THIS USEFUL? 👍 👎

Sophos Access Points

Feedback

1 RELEASE · 2026-08-24

NOTES ↗

Sophos Access Points are wireless networking devices that provide secure network connectivity with integrated threat protection for businesses.

> Sophos AP6 1.10.3173MR10 adds per-SSID traffic shaping, RADIUS-based enterprise authentication (including Microsoft Entra ID), and extends 6 GHz band availability to India and Pakistan.

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

01 **RADIUS server support for AP6 SSIDs** NEW 75

Adds support for configuring RADIUS servers in Sophos Central for AP6 SSIDs with enterprise authentication, including Microsoft Entra ID as a cloud-based identity provider, and adds per-SSID RADIUS server assignment configuration in the local UI.

Enable enterprise Wi-Fi authentication backed by Microsoft Entra ID as a cloud RADIUS identity provider for an AP6 SSID.

> 📍 In Sophos Central, go to RADIUS servers, add a new server selecting Microsoft Entra ID as the identity provider, then assign it to the target AP6 SSID under enterprise authentication settings.

— Covers both Central and local UI RADIUS config with concrete steps.

1.10.3173MR10NewReleased on August 24, …

02 **Traffic shaping for AP6 SSIDs** NEW 65

Adds traffic shaping configuration for AP6 SSIDs under `Advanced settings > Quality of service` in Sophos Central, letting administrators configure bandwidth policies for individual SSIDs.

Configure traffic shaping for a high-priority AP6 SSID to enforce bandwidth limits across the network.

> 📍 In Sophos Central, navigate to the AP6 SSID settings, open Advanced settings › Quality of service, and configure traffic shaping policies for the target SSID.

— Names exact Central UI path and includes step-by-step example.

1.10.3173MR10NewReleased on August 24, …

THINNER COVERAGE BELOW

03 **6 GHz band extended to India and Pakistan** IMPROVED

Extends 6 GHz frequency band availability to India and Pakistan.

— Bare regional expansion statement with no configuration detail.

1.10.3173MR10NewReleased on August 24, …

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Cilium tetragon

Feedback ⓘ 1 RELEASE · 2026-08-25 NOTES ↗

eBPF-based Security Observability and Runtime Enforcement

> Tetragon adds a CEL-expression selector for filtering kernel events in tracing policies (deprecating `MatchArgs`), expands TracingPolicy argument-type and policy-filter scoping, and ships a breaking change to `returnArgAction`.

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

01 CEL expression selector for tracing policies NEW 88

Tracing policies gain a `matchCEL` selector for filtering kernel events using CEL (Common Expression Language) expressions, for example arg1 == int32(-1) && data0 == arg0 + int32(42), with syscall arguments exposed as zero-based indexed variables (`arg0`, `arg1`, ...) and `data:` list items exposed as `data0`, `data1`, ... The older `MatchArgs` operator is deprecated in favor of `MatchCEL`.

Filter a tracing policy to trigger only when `lseek` is called with specific argument relationships, using a CEL expression instead of the deprecated `MatchArgs` operator.

yaml

```yaml
matchCEL:
  expr: "arg1 == int32(-1) && data0 == arg0 + int32(42)"
```

— Concrete syntax, named fields and a runnable YAML example. product docs

02 Deprecation of `returnArgAction: Post` BREAKING 88

`returnArgAction` no longer accepts `Post` as a value in TracingPolicy. Use `returnArg` to include return values in events and omit `returnArgAction` for default behavior; only `TrackSock` and `UntrackSock` remain supported values. Existing policies with `returnArgAction: "Post"` must remove that field.

— Exact migration steps and field names given. v1.7.1

03 Policy filter scope for host, uprobe and USDT events IMPROVED 67

`hostSelector` can now be used in TracingPolicy without requiring the daemon's `enable-policy-filter` flag, reducing configuration overhead for host-scoped policies. Uprobe and USDT events are now scoped through policy filters as well, bringing them in line with existing kprobe/tracepoint filtering.

— Names the flag and event types but no config example. v1.7.1

04 **Expanded argument type support in TracingPolicy** IMPROVED

65

TracingPolicy now supports `btfType` and `btfTypeModule` for more precise BTF-based type resolution of kernel arguments, and adds `int` argument type support for tracepoints, expanding the set of kernel tracepoint arguments that can be traced and filtered.

— Names three type keys but gives no usage example. v1.7.1

THINNER COVERAGE BELOW

05 **Redaction filters applied to all environment variables**

IMPROVED

30

Redaction filters now apply to all environment variables captured in events, broadening the reach of existing redaction configuration.

— Thin description with no config key or mechanism given. v1.7.1

└──▷ **BREAKING ON UPGRADE**

! The `MatchArgs` operator is deprecated in favour of the `MatchCEL` selector.

! `returnArgAction` no longer accepts `Post` as a value in TracingPolicy. Use `returnArg` to include return values in events and omit `returnArgAction` for default behavior; only `TrackSock` and `UntrackSock` remain supported. Existing policies with `returnArgAction: "Post"` must remove that field.

WAS THIS USEFUL?  

Alibaba Cloud OpenSandbox

Feedback ⓘ

1 RELEASE · 2026-08-25 NOTES ↗

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

> OpenSandbox v1.0.19 adds OpenTelemetry tracing for pool warmup, new lifecycle hooks and detached session runs, reworks pool refill and connection pooling internals, and removes the configurable reconciliation interval in a breaking change.

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

01 **OpenTelemetry tracing for pool warmup** NEW 85

Enables OpenTelemetry tracing for pool warmup phases via ConnectionConfig.enableTracing(true), covering queue wait and every warmup stage using the application's global SDK, exporter, sampler, and propagators.

— Names exact config flag and traced stages java/sandbox/v1.0.19

02 **Async pool refill pipeline with rate-limited admission** IMPROVED 85

Stages pool refill into a separate async pipeline covering create, readiness, prepare, post-prepare health checking, renewal, and commit, with per-pool rate-limited create admission via `warmupCreateQps`.

— Names full pipeline stages and rate-limit config key java/sandbox/v1.0.19

03 **Removal of configurable pool reconciliation interval** BREAKING 85

SandboxPool.Builder.reconcileInterval(...) and the corresponding `PoolConfig` setting have been removed; pool reconciliation now runs on a fixed one-second cadence. Callers must delete all uses of `reconcileInterval` with no replacement.

— Exact removed API and required migration action java/sandbox/v1.0.19

04 **Shared OkHttp connection pool in SandboxPool** IMPROVED 80

Adds shared OkHttp connection pool ownership in `SandboxPool`, sized from `warmupConcurrency`, reused across warmup, acquire, and direct-create paths, and evicted on shutdown.

— Names sizing key and lifecycle but no call example java/sandbox/v1.0.19

05 **Detached background runs in persistent sessions** NEW

60

Supports detached background runs inside persistent isolated sessions for Kotlin consumers, with incremental log polling bounded by session lifetime.

– Describes behaviour but no exact API signature `java/sandbox/v1.0.19`

THINNER COVERAGE BELOW

06 **Richer runtime metadata in SDK** NEW

55

Exposes `Sandbox.allocation` summary for confirmed Kubernetes Pool allocations as optional runtime metadata, and models execd hardening capability results as part of the same richer runtime metadata surface in the SDK.

– Names metadata fields but no usage detail `java/sandbox/v1.0.19`

07 **preStart and periodic lifecycle hooks** NEW

50

Adds `preStart` and `periodic` lifecycle hook models accepted at sandbox creation time.

– Names hook types but no config example given `java/sandbox/v1.0.19`

└──▷ **BREAKING ON UPGRADE**

! SandboxPool.Builder.reconcileInterval(...) and the corresponding `PoolConfig` setting have been removed; pool reconciliation now runs on a fixed one-second cadence. Callers must delete all uses of `reconcileInterval` with no replacement.

WAS THIS USEFUL?  

◆ Container & Image Security

Chainguard

Feedback ⓘ 1 RELEASE · 2026-08-24 NOTES ↗

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

> Chainguard's biggest window changes are on the platform side: API v2 reaches general availability with a stable contract and new SDKs, chainctl gains a full set of Catalog Starter management subcommands, and SCIM user provisioning enters Early Access alongside self-serve Helm chart provisioning and 27 new container images.

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

01 Catalog Starter management via chainctl NEW

90

Catalog Starter can now be managed entirely from the CLI: `chainctl starter init` signs up for a new Catalog Starter account interactively, `chainctl starter request-access` requests access to an existing organization's Catalog Starter plan, `chainctl starter add-images` adds up to four images by name to a subscription, and `chainctl starter status` displays registry path, account provisioning status, image quota usage, and per-image readiness.

Sign up for a new Catalog Starter account without leaving the terminal.

```
$ chainctl starter init
```

Add specific images to your Catalog Starter quota after identifying them in the Chainguard Directory.

```
$ chainctl starter add-images python nginx
```

Check quota usage and per-image readiness before pulling images in CI.

```
$ chainctl starter status
```

— All four subcommands named with runnable example commands product docs

02 API v2 general availability IMPROVED

90

Promotes Chainguard API v2 to general availability at `/v2/`, with a stable contract covering organization and access management, registry metadata, vulnerability and

security advisories, event subscriptions, and library artifacts; existing `v2beta1` integrations only need to update the path prefix. The stable API adds cursor-based pagination, server-side ordering, resource naming, dedicated resource lookups, partial updates, and typed error details, plus an OpenAPI specification and v2 Go SDK clients.

— Path, migration note, and SDK/spec named Week of 2026-08-24

03 **SCIM user provisioning (Early Access)** NEW 73

Adds SCIM user provisioning (Early Access) for Okta and Microsoft Entra ID via the `/Users` SCIM resource, creating or updating provisioning records on assignment and blocking login within roughly one hour of deactivation or token expiry.

— Endpoint and IdPs named, no setup steps given Week of 2026-08-24

04 **Self-serve Helm chart provisioning** NEW 60

Enables self-serve Helm chart provisioning from the Chainguard Console for Catalog customers, resolving required images and adding them in a single request, then provisioning asynchronously without EOL tags.

— UI entry point named, no exact steps or command Week of 2026-08-24

THINNER COVERAGE BELOW

05 **27 new container images added** NEW 45

Adds 27 new container images this week including `airbyte-bootloader`, `airbyte-cron`, `airbyte-worker`, `airbyte-workload-api-server`, `airbyte-workload-launcher` (all with FIPS variants), `percona-server-mongodb-fips`, `application-gateway-kubernetes-ingress`, `falcosidekick-ui`, `kubernetes-csi-external-health-monitor-fips`, `mlrun`, `mlrun-api`, `spiffe-csi-driver`, `apache-activemq-artemis-fips`, `aws-iam-controller`, `commercial-nginx-plus-fips`, `kserve-lgbserver`, and `ranger`.

— Names many images but no usage detail Week of 2026-08-24

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ 1 RELEASE · 2026-08-18 BLOG ↗

Anchore's software supply chain solutions automate vulnerability scanning, strengthen container security, and support compliance with NIST, FedRAMP & more.

> Anchore Enterprise extends SBOM generation beyond containers to VM filesystems and disk images, unifying them with container images and third-party SBOMs under one application version, with broadened import/export format support.

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

01 **VM filesystem and disk-image SBOM generation** NEW 95

Adds the `anchorectl app version asset add filesystem` subcommand with a `--type virtual_machine_disk` flag to generate SBOMs from live VM filesystems or mounted disk images and attach them to an application version. Supports offline scanning of VM disk images (read-only, no agent) via QEMU NBD mount, e.g. `sudo qemu-nbd --read-only --connect=/dev/nbd0 disk.qcow2`, enabling forensic and golden-image workflows. Skips pseudo-filesystems (`/proc` , `/sys`) by default when scanning `/` ; narrower paths such as `/usr` , `/opt` , and `/srv` are also supported for faster targeted scans.

Scan a live VM's full filesystem and attach the resulting SBOM to an application version for unified querying alongside container assets.

```
$ anchorectl app add my-vm-fleet --contact-name "Platform Team"
anchorectl app version add 2026-06-01 --app my-vm-fleet
anchorectl app version asset add filesystem / \
  --app my-vm-fleet \
  --version 2026-06-01 \
  --asset web-server-vm \
  --type virtual_machine_disk \
  --wait
```

Generate an SBOM from an offline VM disk image (read-only, no agent required) — useful for forensic analysis or scanning golden-image templates before fleet cloning.

```
$ sudo qemu-nbd --read-only --connect=/dev/nbd0 disk.qcow2
sudo mount -o ro /dev/nbd0p1 /mnt/vm
anchorectl app version asset add filesystem /mnt/vm \
  --app my-vm-fleet \
  --version 2026-06-01 \
  --asset web-server-vm-image \
  --type virtual_machine_disk \
  --wait
```

— Names exact subcommand, flag, and includes runnable commands

What your SBOM inventory misses if it o…

02 **Unified asset model across container, VM and SBOM assets** NEW

70

Unifies container images (`anchorectl app version asset add container-image`), VM filesystems (`anchorectl app version asset add filesystem`), and third-party SBOMs (`anchorectl app version asset add sbom`) as three asset types under one application version, enabling cross-asset querying.

— Names three asset subcommands but lacks a dedicated example

What your SBOM inventory misses if it o…

03 **SBOM format import/export with NTIA minimum elements** IMPROVED

65

Imports and exports CycloneDX (JSON and XML) and SPDX (JSON and tag-value) formats; generated SBOMs include NTIA minimum elements: supplier name, component name, version, `purl` , dependency relationships, SBOM author, and timestamp.

— Names formats and required fields but no usage command

What your SBOM inventory misses if it o…

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Sumo Logic

Feedback ⓘ 1 RELEASE · 2026-08-25 BLOG ↗

Sumo Logic is a cloud-based platform for real-time log management, monitoring, and security analytics across applications and infrastructure.

> Sumo Logic dashboards can now query beyond the previous 32-day time range limit when backed by aggregate scheduled searches or views.

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

01 **Extended dashboard time ranges beyond 32 days** IMPROVED

45

Dashboards now support time ranges beyond the previous 32-day limit when backed by aggregate scheduled views, aggregate scheduled searches, or a combination of both.

— Names the mechanism and prior limit but no config steps.

Support for Longer Dashboard Time Range…

WAS THIS USEFUL? 👍 👎

Google SecOps

Feedback

ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

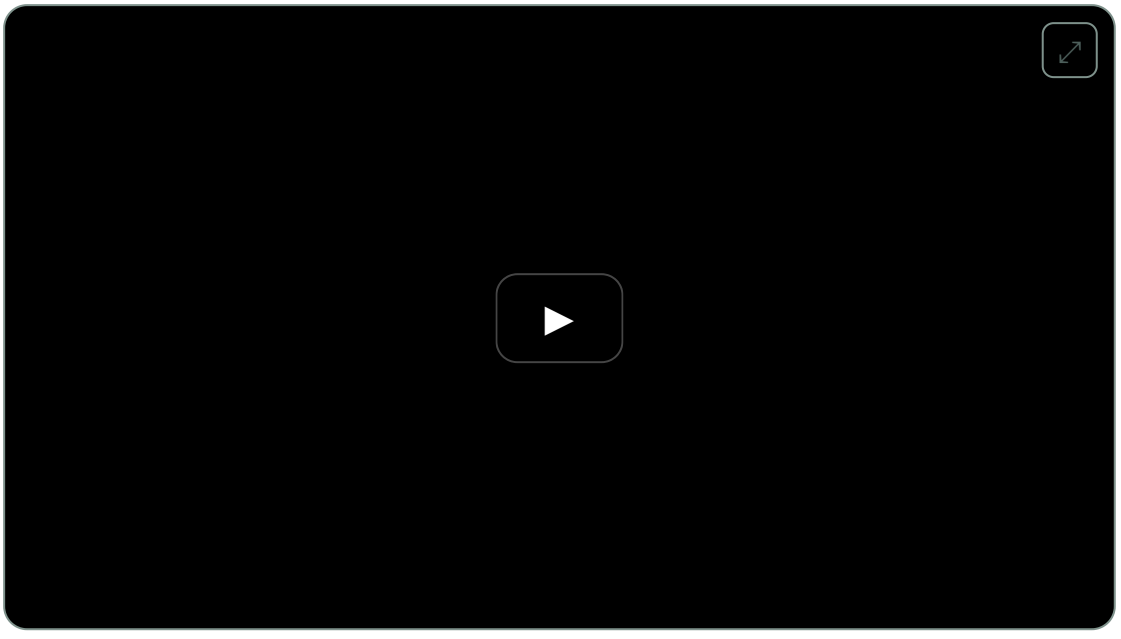
> Google SecOps shipped a side-by-side alert inspection layout in Cases, precise relative time filter operators for dashboards and Search, localized Operations intel in the Emerging Threats Center, and a new Unroll processor for splitting array-based log entries in data pipelines.

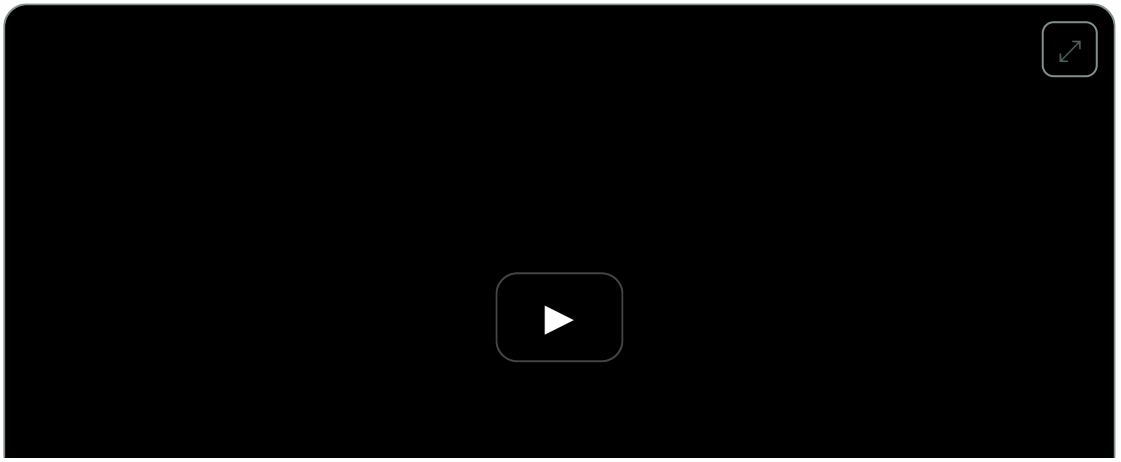
└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

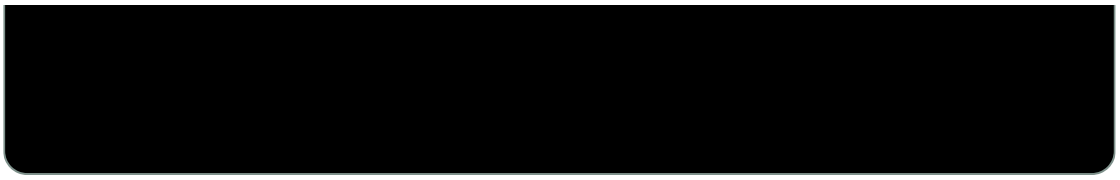
01 Unroll processor for array log events in pipelines NEW

88

A new Unroll processor (event breaking) in Google SecOps data processing pipelines automatically expands log entries containing arrays or slices into multiple individual log events prior to parsing and ingestion. It requires structured input, so raw string payloads must first be converted using a Transform processor (e.g. set(body, ParseJSON(body))) placed before the Unroll processor in the pipeline execution sequence.







Split JSON array log payloads into individual events in a pipeline — place the Transform processor before the Unroll processor so raw strings are parsed first.

```yaml
# Pipeline processor sequence (data processing pipeline config)
processors:
  - type: Transform
    expression: set(body, ParseJSON(body))
  - type: Unroll
```

— Names exact processors, function, and runnable pipeline config

August 24, 2026 (2026-08-24)

02 **Relative time filter operators in dashboards and Search** NEW

65

Three new relative time filter operators — Past, Previous, and Current — are available in SecOps dashboards and Search to eliminate ambiguity between rolling windows and calendar-aligned periods, working across all time units from seconds through years. Currently in public preview.

Use the new Previous operator to filter dashboards to a complete prior calendar period (e.g. last full week) instead of a rolling window.

📍 In Google SecOps Search or a dashboard, open the time filter picker › select Previous › choose the time unit (e.g. Week) to scope results to the last complete calendar week.

— Named operators and UI steps but no config surface product docs

THINNER COVERAGE BELOW

03 **Side-by-side alert inspection in Cases** NEW

58

A new Side-by-side view layout on the Alerts & Detections tab in Cases lets analysts inspect alert metadata, status, priority, creation date, and Gemini investigation insights in an adjacent side pane without leaving the main alert list.

Switch to side-by-side alert inspection in Cases to review Gemini insights and metadata without losing your place in the alert list.

📍 In Google SecOps, go to Cases › open a case › Alerts & Detections tab › click Side-by-side view in the layout toggle.

— UI path given but no deeper mechanism detail product docs

04 **Operations added to Emerging Threats Center** NEW 38

The Emerging Threats Center feed now includes Operations, surfacing localized threat intelligence from frontline investigations (such as Managed Threat Defense engagements) alongside global Campaigns to map full adversary TTPs.

— Explains what it does but no usage steps given product docs

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-08-24 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI's latest release focuses on scaling stream infrastructure by offloading oversized Redis events to file storage/S3, alongside new scoring fields for Incidents and Events, faster STIX bundle processing, experimental cross-platform Status sync, and a streamlined One-Pager Fintel report mode.

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

01 Large Redis stream events offloaded to file storage NEW

86

A new `redis:max_event_length` config key (env: `REDIS__MAX_EVENT_LENGTH`) sets a maximum size in bytes for a Redis stream event; content exceeding the limit is offloaded to file storage automatically, and setting it to `0` disables the cap. This was extended so large stream event messages are stored to S3 instead of Redis, removing Redis size constraints for high-volume stream deployments.

Cap Redis stream events at 1 MB to prevent oversized payloads from bloating the stream — content exceeding the limit is offloaded to file storage automatically.

```yaml
redis:
  max_event_length: 1048576
```

— Names config key, env var, default disable value, and S3 mechanism 7.260824.0

THINNER COVERAGE BELOW

02 One-Pager mode for Fintel reports NEW

45

Introduces One-Pager mode for Fintel reports, making the first and back cover pages optional for single-page intelligence products.

— Describes behaviour but no navigation or config specifics 7.260824.0

03 Node.js STIX bundle splitting before worker queue IMPROVED

43

Splits STIX bundles in Node.js before pushing to the worker queue, reducing worker-side processing overhead for large bundle imports.

— Explains mechanism but no config or measurable numbers given 7.260824.0

04 **Score field on Incident and Event objects** NEW 43

Adds a `score` field to the Incident data model and a `score` field to the Event data model, enabling confidence/severity scoring on both object types.

— Names exact field and two models but no usage mechanism 7.260824.0

05 **Experimental cross-platform Status sync** NEW 35

Adds experimental Status sync between OpenCTI platforms via live streams, laying groundwork for cross-platform workflow state propagation.

— Mechanism named but marked experimental with no usage detail 7.260824.0

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

Tenzir is a data pipeline platform that ingests, transforms, and analyzes security telemetry and logs for threat detection.

> Tenzir Ship v2.2.0 introduces forward-looking metadata enforcement for changelog entries, adding required-author and no-PR policies plus a way to audit historical releases against them.

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

01 Changelog metadata enforcement policies in Tenzir Ship

95

BREAKING

Changelog project configuration gains `require_author: true` to require at least one author on unreleased entries and `omit_pr: true` to reject pull request numbers in unreleased entries. The `validate` command now applies the `omit_pr` and `omit_author` policies only to unreleased entries by default; pass `--all-entries` to audit every past release against the current policy instead. This is a behavior change: `validate` no longer checks released/published entries against `omit_pr` / `omit_author` by default, so workflows relying on full-history omission checks must add `--all-entries` to restore prior behaviour.

Enforce author attribution on all unreleased changelog entries in CI.

yaml

```yaml
require_author: true
omit_pr: true
```

Audit every historical release against the current metadata policy, not just unreleased entries.

```
$ tenzir-ship validate --all-entries
```

— Names exact config keys, validate flag, and the default-behavior change.

Tenzir Ship v2.2.0: Forward-looking cha…

└──▷ BREAKING ON UPGRADE

! The `validate` command no longer applies `omit_pr` and `omit_author` policies to released (published) entries by default; existing workflows that relied on full-history omission checks must now pass `--all-entries` to restore that behaviour.

WAS THIS USEFUL?

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ 1 RELEASE · 2026-08-25 CHANGELOG ↗

runZero is a cloud-based asset discovery and management platform that identifies and inventories all connected devices across networks.

> runZero added OCSF 1.9.0 export support for asset inventories and improved asset count visibility in organization overview reports.

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

01 **OCSF 1.9.0 export for asset inventory** NEW 83

runZero now supports exporting the asset inventory in OCSF 1.9.0 format, producing `assets.ocsf.jsonl` and `assets.ocsf.json` files accessible from the inventory export menu and the Export API. Assets are exported as Device Inventory Info events and observed services as Live Evidence Info events, with any unmapped fields preserved under the event's `unmapped` object.

— Names exact file formats, event types and unmapped field handling.

changelog-20260825-5536acfb

THINNER COVERAGE BELOW

02 **Asset counts in organization overview report** IMPROVED 35

The organization overview report now shows total asset counts in the sites summary, matching the sites table already shown in the runZero Console.

— Describes the change but no config or navigation detail. changelog-20260825-5536acfb

WAS THIS USEFUL?

Intruder

Feedback

ⓘ

1 RELEASE · 2026-08-25

BLOG ↗

Give feedback to the Intruder team so we can make more informed product decisions. Powered by Canny.

> Intruder's AI Pentesting now supports Bitbucket Cloud repositories, joining GitHub and GitLab as source-control integrations.

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

01 **Bitbucket Cloud integration for AI Pentesting** NEW

55

Adds Bitbucket Cloud as a supported source-control integration for AI Pentesting, alongside existing GitHub and GitLab support. Connect repositories via the Integrations page in the Intruder console to run AI-powered penetration tests against Bitbucket-hosted code.

Enable Bitbucket Cloud so your Bitbucket-hosted repositories appear as targets for AI Pentest scans.

> 📍 In the Intruder console, go to Integrations, find Bitbucket Cloud, and enable the integration to connect your repositories.

— Names integration and UI path but no technical mechanism.

Integrate with Bitbucket Cloud for AI P…

WAS THIS USEFUL? 👍 👎

Censys

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

Censys is a search engine for discovering internet-connected devices, servers, and certificates to assess security exposure and vulnerabilities.

> Censys expanded its Platform Audit Log with 35 new event types spanning ASM, SAML, and user management, and shipped a geo_distance CenQL function, expanded HTTP body search, 14 new protocol/application scanners, and 2 new software fingerprints.

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

01 **geo_distance function in CenQL** NEW 90

Adds `geo_distance` function to Censys Query Language (CenQL) for finding assets relative to a geographic location; it currently supports `host.location.coordinates` and `host.services.ntrip` as target data objects and is available to all users.

Find exposed NRPE agents (Nagios Remote Plugin Executor) within 50 km of a specific data center to assess monitoring infrastructure attack surface.

```
$ services.service_name: NRPE AND
geo_distance(host.location.coordinates, 37.7749, -122.4194, 50km)
```

Hunt for KNX/IP building-automation endpoints near a target facility — useful for OT/ICS exposure assessments.

```
$ services.service_name: KNXIP AND
geo_distance(host.location.coordinates, 51.5074, -0.1278, 25km)
```

— Named function, supported fields, and runnable query examples

August 24, 2026 (2026-08-24)

02 **Expanded HTTP body search size** IMPROVED 80

Expands searchable HTTP body size for `web.endpoints.http.body` (name-based web property endpoints) from 2 KB to 6 KB, enabling more targeted Platform queries.

Search HTTP bodies of web endpoints for a credential string or internal banner that only appears deeper than 2 KB, now reachable with the expanded 6 KB body search.

```
$ web.endpoints.http.body: "Authorization: Bearer internal-" AND
web.endpoints.http.body: "env=production"
```

— Exact field, before/after limits, and a runnable query August 24, 2026 (2026-08-24)

03 **Expanded Platform Audit Log event coverage** NEW 70

Platform Audit Log adds 35 new event types across five areas: ASM risk management (ASM Exclude Created , ASM Exclude Deleted , ASM Risk Instance Accepted , ASM Risk Instance Unaccepted , ASM Risk Instance Severity Changed , ASM Risk Type Disabled , ASM Risk Type Enabled , ASM Risk Type Severity Changed , ASM Seed Created , ASM Seed Deleted), SAML configuration lifecycle (SAML Config Created , SAML Config Deleted , SAML Config Domain Verified , SAML Config Updated , available on Platform and ASM), Personal Access Tokens (PAT Created , PAT Deleted , Platform only), user authentication and account lifecycle (User Login , User Login Failed , User MFA Changed , User Password Changed , User Password Reset , User Settings Changed , User Created , User Disabled , User Updated), and organization/membership management (Organization Created , Organization Updated , Organization Deleted , Membership Created , Membership Removed , Membership Updated , Invitation Accepted , Invitation Created , Invitation Deleted , Invitation Resent).

.
.
.

— Names every event but log review is passive, not a direct action product docs

04 **14 new protocol and application scanners** NEW 70

Adds scanners for `MPLS_LDP` (MPLS Label Distribution Protocol, RFC 5036, all users), `SMUX` (SNMP Multiplexing Protocol, RFC 1227, all users), `IPERF3` (iPerf3 bandwidth-measurement control protocol, all users), `TVT_CONTROL` (TVT proprietary remote-management/CMS protocol for white-label DVR/NVR/IP camera firmware, all users), `DLT` (AUTOSAR Diagnostic Log and Trace / dlt-daemon for automotive ECU logging, Critical Infrastructure module users), `DUBBO` (Apache Dubbo RPC framework, Starter/Search/Core users), `COMELIT_ICONA` (Comelit ICONA Bridge intercom protocol, all users), `FOURD_SERVER` (4D Server client-server publish protocol, Starter/Search/Core users), `KNXIP` (KNXnet/IP building-automation bus, Starter/Search/Core users), `NRPE` (Nagios Remote Plugin Executor agent, all users), `REDSHIFT` (AWS Redshift analytics DB, Starter/Search/Core users), `MANTICORE_SPHINXAPI` (Manticore Search / Sphinx legacy SphinxAPI protocol, Starter/Search/Core users), `HOOSAT_HTND_GRPC` (Hoosat Network htnd full-node gRPC API, all users), and `TAILSCALE_DERP` (Tailscale Designated Encrypted Relay for Packets, Starter/Search/Core users).

— Every scanner and tier named, but no query example attached

August 24, 2026 (2026-08-24)

THINNER COVERAGE BELOW

05 **New software fingerprints for Bifrost and Juniper Junos OS** NEW 50

Adds fingerprint for Bifrost software (AI gateway unifying access to OpenAI, Anthropic, AWS Bedrock, and Google Vertex via an OpenAI-compatible API) and fingerprint for Juniper Junos OS (modular network OS for Juniper routing, switching, and security devices).

— Names both fingerprints but no query or usage shown August 24, 2026 (2026-08-24)

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo

Feedback ⓘ 1 RELEASE · 2026-08-24 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo Pro's biggest additions this window are three enterprise deployment options — an air-gapped install mode, official OpenShift 4.x support, and FIPS 140-3 validated images — plus a new interactive Asset Hierarchy diagram, tighter product-scoping on endpoint operations, mandatory verification for MFA disable, and a wide documentation expansion.

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

01 **Air-gapped install mode in dojo-compose-cli** NEW 95

Enabled via `dojo-compose-cli config set --air-gapped true`, which validates the license from `/etc/defectdojo/dojopro.lic` and blocks registry-dependent commands with an explicit error instead of failing silently. Supporting commands: `--deploy-version` and `--version` flags to pin the exact release across artifacts and images; `dojo-compose-cli deploy download` to pull the full deployment directory into `/opt/dojo`; `dojo-compose-cli app pull-images` to pre-fetch container images on a staging host; `dojo-compose-cli environment print` to capture generated secrets (`DD_CREDENTIAL_AES_256_KEY`, `DD_SECRET_KEY`); `dojo-compose-cli environment add --key ... --value ...` to set values (`DD_CREDENTIAL_AES_256_KEY`, `DD_SECRET_KEY`, `DD_SITE_URL`, `DD_ALLOWED_HOSTS`, `DD_DATABASE_URL`) on the air-gapped target; and `dojo-compose-cli app start` / `stop` / `restart`, which operate fully offline against the local compose file. Requires CLI 2.1.0+ — earlier versions have no air-gapped mode and fail with name resolution errors. The `DOJO_CLI_KEY` environment variable supplies the CLI's config encryption key non-interactively.

Enable air-gapped mode on the target host so the CLI reads the license from disk and refuses registry calls rather than hanging or failing silently.

```
$ sudo ./dojo-compose-cli config set --air-gapped true
dojo-compose-cli config print
```

On the staging host, set the target release and pre-pull all images before packaging them for transfer into the air-gapped network.

```
$ export DOJO_CLI_KEY="your-key"
dojo-compose-cli config set --deploy-version 2.1.0
dojo-compose-cli config set --version 2.1.0
dojo-compose-cli deploy download
dojo-compose-cli app pull-images
docker image ls
```

Inject required secrets and database connection into the air-gapped host's configuration after restoring the deployment directory.

```
$ dojo-compose-cli environment add --key
"DD_CREDENTIAL_AES_256_KEY" --value "<64-character-key-from-step-
4>"
dojo-compose-cli environment add --key "DD_SECRET_KEY" --value "
<64-character-key-from-step-4>"
dojo-compose-cli environment add --key "DD_DATABASE_URL" --value
"postgres://<db_user>:<db_password>@<db_host>:5432/<db_name>"
dojo-compose-cli environment add --key "DD_SITE_URL" --value
"https://defectdojo.internal.example.com"
dojo-compose-cli environment add --key "DD_ALLOWED_HOSTS" --value
"*"
dojo-compose-cli app start
```

— Full runnable CLI command set with exact flags and env vars `product docs`

---

02 **Official OpenShift 4.x deployment support** NEW 90

Adds config keys `dojo.secureCookies` (must be `false` behind an edge-terminated Route or login fails), `dojo.skipInitContainer` (set `true` after first install to skip the initializer wait on upgrades), and `config.logLevel` / `celery.logLevel` (accept `DEBUG`). Supports three resource profiles — `minimal`, `standard`, `performance` — the last enabling autoscaling, and ships an idempotent bootstrap script with a dry-run mode that handles storage, `fsGroup`, and Route creation. Runs under the default `restricted-v2` SCC on OpenShift Container Platform, ROSA, and OKD with no `anyuid` / `privileged` grant required, exposing the service via an edge-TLS Route with HTTP-to-HTTPS redirect (ROSA hostnames auto-generated as `<release-name>-<namespace>.apps.<cluster-domain>`). Requires ReadWriteMany storage for multi-node deployments, with NFS-backed `fsGroup` derived from the namespace's `openshift.io/sa.scc.supplemental-groups` annotation; exploitability enrichment can optionally pull EPSS scores from `api.first.org` and CISA KEV data from `www.cisa.gov` over HTTPS.

Fix login failures behind an edge-terminated OpenShift Route by disabling secure cookies in your Helm values.

```yaml
dojo:
  secureCookies: false
```

Retrieve the correct fsGroup start value from the namespace annotation when using NFS or EFS storage.

```
$ oc get namespace <namespace> -o
jsonpath='{.metadata.annotations.openshift\.io/sa\.scc\.supplemen
tal-groups}'
```

Skip the initializer wait on subsequent upgrades once the database has already been initialized.

```yaml
dojo:
  skipInitContainer: true
```

— Named Helm keys and exact SCC/Route/storage mechanics `product docs`

03 **FIPS 140-3 validated container images** NEW 90

Adds an `-fips` image tag suffix for FIPS 140-3 validated variants (standard images unchanged) and a `DD_FIPS_MODE="1"` environment variable to activate enforcement across Docker Compose, Kubernetes, and ECS. Helm values `fips.enabled` auto-select FIPS images and inject `DD_FIPS_MODE` chart-wide, while `fips.validate: false` explicitly accepts partial coverage when non-FIPS components like Sensei or embedded datastores are also enabled; the chart otherwise refuses to render and names the offending components. Cryptographic operations run through OpenSSL FIPS Provider 3.1.2 (CMVP #4985) for Python/Django services and Go Cryptographic Module v1.0.0 (CMVP #5247) for Go services. FIPS mode works without a FIPS-enabled host kernel, including on ECS Fargate, with documented task definitions using `DD_FIPS_MODE`, `DD_DATABASE_URL`, `DD_SECRET_KEY`, `DD_CREDENTIAL_AES_256_KEY`, and `DD_LICENSE` sourced from Secrets Manager.

Enable FIPS 140-3 mode in a Docker Compose deployment by pointing to FIPS image variants and setting `DD_FIPS_MODE` in each shared environment anchor.

```yaml
# In .env
DD_IMAGE_TAG=<version>-fips

# In docker-compose.yaml shared anchors
x-dojo-vars: &dojoenv
  DD_FIPS_MODE: "1"

x-nginx-vars: &nginxenv
  DD_FIPS_MODE: "1"

x-psirt-vars: &psirtenv
  DD_FIPS_MODE: "1"

# Then recreate the stack
docker compose up -d --force-recreate
```

Deliberately include Sensei in a FIPS environment while documenting the bundled scanner toolchain as non-validated in the system security plan.

```yaml
fips:
  enabled: true
  validate: false
sensei:
  enabled: true
```

— Named env vars, Helm keys and specific CMVP certificate numbers `product docs`

04 **Interactive Asset Hierarchy diagram** NEW 83

Visualizes parent-child Asset relationships as labelled, color-coded arrows with a `Legend` panel, and a `Hide Asset List` button to give the diagram full-page space. Per-node display can toggle `Asset ID`, `Organization ID`, `Organization Name`, and `Child Count`. Clicking a node exposes `Open Asset`, `Edit Asset`, `Add Child`, `Change Parent`, and `Remove From Hierarchy` actions without leaving the diagram: `Add Child` nests an existing or new Asset within the same Organization, `Change Parent` moves an Asset to a different parent, and `Remove From Hierarchy` detaches it with a prompt to orphan or re-parent its children. Large hierarchies load incrementally via `Load Parents` and `Load` (with configurable batch size).

Reorganize a large product tree without losing context — move an Asset to a new parent while keeping the diagram visible.

📍 In the Asset Hierarchy diagram, click the Asset node you want to move. In the right panel, click `Change Parent` and select the new parent Asset from the form that opens in the right panel.

Explore a deep hierarchy incrementally to avoid an overwhelming diagram — load ancestors of an unparented node on demand.

📍 In the Asset Hierarchy diagram, locate an Asset whose parent has not been loaded and click the `Load Parents` button that appears above it to add the parent and its siblings.

— Every UI control named, but only navigation not runnable commands `product docs`

THINNER COVERAGE BELOW

05 **Documentation site infrastructure improvements** NEW

35

Adds an RFC 9116 `security.txt` to `docs.defectdojo.com`, adds Brazilian Portuguese, Simplified Chinese, and Italian translations of the core documentation guides, and splits connector documentation into one page per connector for easier navigation.

— Docs-site changes only, no product behavior described `3.2.300`

06 **Documentation expansion across product areas** IMPROVED

35

Newly documented, previously-existing capabilities include typed Custom Fields and their integration with Rules Engine 2.0, multi-Asset finding fan-out in Smart Upload workflows, AI Model Settings and supported LLM providers in Sensei, GMP over SSH support for OpenVAS/Greenbone connectors, customizable page layouts on view pages and Form Configuration in navigation, a Quick Export dialog that produces a named report, the rebuilt Asset hierarchy page, and a four-permission rule tier in Rules Engine 2.0.

— Each item is a one-line doc mention with no mechanism `3.2.300`

07 **Endpoint delete/rename scoped to product** BREAKING

30

Endpoint delete and rename operations are now scoped to the requesting product, restricting cross-product mutations that were previously possible.

— No endpoint name or prior behavior detail given `3.2.300`

08 **MFA disable requires verification code** IMPROVED

28

Disabling multi-factor authentication now requires entering a verification code, adding a confirmation step before MFA can be turned off.

— Brief behavioral note without flow or endpoint detail `3.2.300`

WAS THIS USEFUL?  

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

iDenfy

Feedback ⓘ 1 RELEASE · 2026-08-25 CHANGELOG ↗

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy's Sprint 133 release adds a new API endpoint to re-run KYB automation, expands identity document coverage with Brazil's Digital CNH across web and mobile SDKs, and introduces smaller dashboard and branding tweaks alongside brief mentions of standalone bank card verification, Persian/Farsi RTL support, and a Risk Assessment PDF export.

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

01 **Digital CNH document support across platforms** NEW

55

Brazil's Digital CNH (Carteira Nacional de Habilitação) is now supported in identity verification. Android SDK 9.1.1 and iOS SDK 9.1.1 were both updated to add Digital CNH scanning support, and the Android SDK update also improved virtual camera detection.

— Names document type and exact SDK versions but no usage steps

changelog-20260825-26c93dcc

02 **KYB automation re-run API endpoint** NEW

45

A new API endpoint lets teams programmatically retrigger KYB automation workflows, allowing KYB checks to be re-run without manual intervention.

— Endpoint named but no path, method, or parameters given changelog-20260825-26c93dcc

03 **Persian/Farsi RTL support** NEW

10

Right-to-left (RTL) language support for Persian/Farsi was added, per the release summary.

— Only named in summary, no further detail provided changelog-20260825-26c93dcc

WAS THIS USEFUL? 👍 👎

Clerk

Feedback

1 RELEASE · 2026-08-25

CHANGELOG ↗

Clerk is an authentication and user management platform that handles sign-up, login, and identity verification for applications.

> Clerk shipped custom OAuth scope management for MCP clients and biometric trusted-device sign-in across its mobile SDKs (Expo, iOS, Android) with matching prebuilt UI support.

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

01 Trusted-device biometric sign-in across SDKs NEW 90

Adds device-bound biometric authentication (Face ID, Touch ID, Android biometrics) as a passwordless sign-in option, implemented per platform: `useTrustedDevices` hook with `enroll`, `signIn`, and `revoke` methods in `@clerk/expo`; `Clerk.shared.trustedDevices.enroll`, `Clerk.shared.auth.signInWithTrustedDevice`, and Clerk.shared.trustedDevices.revoke(id: trustedDeviceId) on iOS; and `Clerk.trustedDevices.enroll`, `signInWithTrustedDevice`, and Clerk.trustedDevices.revoke(trustedDeviceId) on Android. Prebuilt auth and user profile views now surface a biometric enrollment prompt, a biometric sign-in button, and a current-device toggle without custom code.

Enroll the current device for biometric sign-in in an Expo app so returning users can authenticate with Face ID or Touch ID instead of a password.

javascript

```javascript
import { useTrustedDevices } from '@clerk/expo';

const { enroll } = useTrustedDevices();
await enroll();
```

Authenticate a returning user on iOS using a device-bound biometric challenge instead of credentials.

swift

```swift
await Clerk.shared.auth.signInWithTrustedDevice()
```

— Names every API method per platform with runnable code examples

changelog-20260825-1aeedc1e

02 Custom OAuth scopes for MCP clients NEW 70

Adds custom OAuth scope definitions in the Clerk Dashboard (examples: `messages:read`, `tools:execute`, `resources/files:read`, `mcp_all`) so MCP clients can request only the access each OAuth application needs. Also adds OAuth metadata advertisement control so MCP clients can discover which scopes an application supports via Clerk's OAuth metadata endpoint.

Define and assign custom OAuth scopes for an MCP client in the Clerk Dashboard so it can only request the permissions your API exposes.

> 📍 In the Clerk Dashboard, go to `Scopes` and create custom scopes (e.g. `messages:read`, `tools:execute`, `resources/files:read`, `mcp_all`). Then open your OAuth application settings and assign only the scopes it requires. Under OAuth metadata settings, choose which scopes to advertise so MCP clients can discover them.

— Names dashboard flow and example scopes but not the metadata endpoint path

`changelog-20260825-1aeedc1e`

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ 1 RELEASE · 2026-08-24 BLOG ↗

SailPoint is an identity governance platform that manages user access, permissions, and compliance across enterprise applications and systems.

> SailPoint's ISC Workflows HTTP Request action can now mask sensitive header values so secrets no longer appear in plain text.

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

01 Sensitive Request Headers in HTTP Request action NEW

69

The HTTP Request action in Identity Security Cloud Workflows adds a Sensitive Request Headers field — a dedicated masked key-value store for API keys, bearer tokens, and other secrets that is encrypted in workflow definitions and omitted from execution history. Users move secret values (e.g. an Authorization key with a `Bearer <token>` value) from the plain-text Request Headers field into this new field when editing the action.

Move secret values out of plain-text Request Headers into Sensitive Request Headers when editing an HTTP Request action, so tokens no longer appear in workflow definitions or execution history.

📍 In a Workflow, open an HTTP Request action › locate the Sensitive Request Headers field › add your key (e.g. Authorization) and secret value (e.g. `Bearer <token>`) there instead of in Request Headers › save the workflow.

— Names exact field and shows step-by-step usage, but scope is narrow.

Enhancement: HTTP Request Action - Sens…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 25, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue