THE CYBER TOOLCHAIN

NO. 050

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 19, 2026 · EVERY DAY

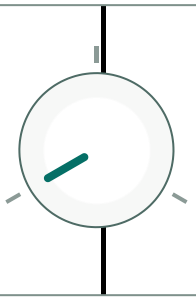
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 50 tools

OFFENSIVE SECURITY

vim ? help

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ 1 RELEASE · SEEN 2026-08-19 NOTES ↗

> Dreadnode's biggest change this window is a versioned item type registry for structured agent output, alongside experimental ScopeGuard policy support in remote sessions, a full Caido MCP server, and a new web-security skill for response-queue poisoning, plus a batch of smaller UI and CLI polish.

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

01 **Versioned item type registry for agent output** NEW

75

Platform admins can define, publish, and manage structured output contracts (findings, assets, and custom types) via Operations → Objects, with capabilities binding to approved registry versions and all item writes validated against the canonical schema.

— Names exact UI path and validation mechanism, no example given

Versioned item type registry for struct…

02 **Caido MCP server in web-security capability** NEW 75

Adds the full Caido MCP server (66 tools, 6 resources) as a second MCP server in the web-security capability, alongside the existing Python wrapper, both targeting the same Caido instance.

— Concrete tool/resource counts and integration scope given

Versioned item type registry for struct…

03 **`--include-public` flag documentation for Hub browse** IMPROVED

60

Adds a `--include-public` flag documentation clarification for Hub browse pages (Datasets, Models, Capabilities, Environments), distinguishing CLI `--include-public` behavior from the Public scope tab.

— Named flag and pages, purely a doc clarification

Versioned item type registry for struct…

THINNER COVERAGE BELOW

04 **Response-queue-poisoning skill for web-security** NEW

55

Adds a `response-queue-poisoning` skill to the web-security capability, covering CRLF-powered desync attacks including response queue poisoning and request tunnelling via HTTP header injection.

– Names attack technique but no usage walkthrough

Versioned item type registry for struct…

05 **Progress indicator for dn evaluation create** IMPROVED 55

Adds a spinner (or plain status line in non-interactive output) to `dn evaluation create` while the evaluation is being submitted.

– Exact command named, behavior is minor UX only Versioned item type registry for struct…

06 **Sandbox isolation reference documentation** NEW 55

Adds a new sandbox isolation reference page covering Firecracker compute isolation, application authentication, public sandbox routing, outbound internet access, credential scope, and shared responsibility.

– Lists topics covered but no page location or link

Versioned item type registry for struct…

07 **ScopeGuard policy support in remote sessions** NEW 50

Adds `ScopeGuard` policy support in remote and headless sessions via the SDK (experimental); the session UI renders tool-call policy decisions when ScopeGuard metadata is present — opt-in, with no workspace setting or picker entry yet.

– Named but experimental with no SDK usage shown Versioned item type registry for struct…

08 **Bulk session delete in agent sessions UI** NEW 45

Multiple agent sessions can now be selected and deleted together, with deleted sessions disappearing from the list immediately without a page refresh.

– UI behavior described but no exact navigation path

Versioned item type registry for struct…

09 **Capabilities UI redesign** IMPROVED 45

Redesigns the capabilities UI with improved navigation, richer detail views for agents, skills, and MCP servers, and a new file browser for capability contents.

– General redesign described without specific navigation steps

Versioned item type registry for struct…

10 **Model slug shown in session metadata** IMPROVED 35

Makes the model slug visible by default in session metadata so the exact model used is identifiable without extra clicks.

– Simple UI visibility change, thinly described Versioned item type registry for struct…

11 **Reader-role org restriction in workspace creation** IMPROVED

Reader-role orgs are now visibly marked and disabled in the New Workspace dialog to prevent 403 errors on form submission.

— Names dialog and error prevented, minor fix Versioned item type registry for struct…

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

Nmap

Feedback ⓘ changes since 2026-08-09 CODE ↗

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

Nmap gains a libfuzzer-based fuzzing framework covering Ncat HTTP, traceroute, DNS, and packet validation routines.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-09
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-09
```

Adds a libfuzzer-based fuzzing framework for Nmap internals, with fuzzers for Ncat HTTP functions, decode_reply() in traceroute, `gettcpopt_ts`, validatepkt(), and several libnetutil functions.

Switches DNS transaction IDs (`TXID`) from sequential to random to reduce predictability in DNS requests.

Improves DNS resolver correctness: timeout now starts only after the write completes, late replies are accepted without clobbering existing results, and request ID lookup is tracked per-server.

Improves Lua's random number generator to produce more random output.

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

SpecterOps BloodHound

Feedback ⓘ 1 RELEASE · 2026-08-18

NOTES ↗

Six Degrees of Domain Admin

> BloodHound v9.6.0 introduces GitHub Actions self-hosted runner attack path modeling alongside a wave of new management, artifact and findings APIs, plus PostgreSQL storage and OpenGraph improvements.

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

01 Client artifact management and S3-backed storage APIs NEW

90

Adds a download client artifact endpoint (`BED-8091`) and a delete client artifact endpoint (`BED-8129`) for retrieving and removing collector client artifacts via the API. Adds `part_number` support to the artifact upload API for multipart upload handling, S3 configuration support (`BED-9174`) to enable object-storage-backed artifact handling, and exposes artifact metadata on management operation responses.

— Names exact endpoint IDs and config field v9.6.0

02 Kind Metadata APIs for node and relationship kinds NEW

85

Adds Kind Metadata APIs to retrieve node and relationship kind details by graph-assigned kind ID, enabling OpenGraph automation and integrations. Specifically adds the `api/v2/relationship-kinds/{relationship_kind_id}` endpoint to retrieve a specific relationship kind by ID, and a new GET node kind endpoint (`BED-8756`) for querying individual node kinds via the API.

— Names exact endpoint paths and internal ticket IDs v9.6.0

03 PostgreSQL graph storage and deletion optimization IMPROVED

70

PostgreSQL Graph Storage Optimization runs after ingest and analysis to improve reliability and performance of PostgreSQL-backed graph databases, including faster graph data deletion. This includes optimized full graph database wipes for faster data-clearing operations in DB Management, and optimized `source_kind` and relation deletion using database indices for improved performance.

— Names mechanism and indices but no exposed config v9.6.0

04 OpenGraph extension management enabled by default IMPROVED

60

Extension Management is now enabled by default — OpenGraph extensions can be managed without manually enabling the feature first, via the `opengraph_extension_management` feature flag being turned on for all users.

— Names the feature flag but no config steps v9.6.0

THINNER COVERAGE BELOW

05 **GitHub self-hosted runner attack path modeling** NEW

55

Adds GitHub Self-Hosted Runner Modeling to identify membership, delegation, and access paths for enterprise and organization-scoped GitHub Actions self-hosted runners and runner groups.

— Describes scope but no UI path or query given product docs

06 **Management operation queue APIs** NEW

50

Adds queue management operation endpoints for submitting and managing background operations via the API, plus a get-available management operation endpoint to discover which operations can be queued.

— No exact endpoint paths given v9.6.0

07 **Client support bundle download and removal** NEW

50

Adds a download client support bundle capability (API endpoint and UI) for pulling diagnostic bundles from managed clients, and a delete client support bundle UI action for removing previously downloaded diagnostic bundles.

— UI actions named but no exact endpoint paths v9.6.0

08 **Findings Table severity filter and finding types endpoint** NEW

50

Adds a findings table severity filter, enabling practitioners to narrow the Findings Table by severity level, and a new Finding Types endpoint for the Findings Table to retrieve available finding type classifications.

— UI filter and endpoint named without exact path v9.6.0

09 **OpenGraph Markdown schema validation** NEW

45

Adds OpenGraph Markdown Validation so extension definition schemas are checked for safe rendering of custom Markdown content before upload.

— No config key or command given product docs

10 **Alert event publisher** NEW

40

Adds an alert event publisher (`BED-7948`) for emitting events when alert conditions are triggered.

— Names ticket ID but not endpoint or event format v9.6.0

11 **Findings Prioritization feature flag** NEW 35

Adds a Findings Prioritization feature flag (`BED-9035`) to gate prioritization workflows.

— Flag named but workflow itself not described v9.6.0

12 **Webhooks table wired up in UI** NEW 30

Wires up the webhooks table in the UI, surfacing configured webhook endpoints.

— Bare UI mention without further detail v9.6.0

13 **Collector guidance for case-sensitive OpenGraph Object IDs**

IMPROVED 25

Adds collector compatibility guidance for case-sensitive OpenGraph Object ID ingest.

— Bare docs note without specifics product docs

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-08-19

NOTES ↗

Horizon3 NodeZero is an autonomous security tool that continuously identifies and exploits vulnerabilities in networks to assess organizational risk.

> Horizon3 NodeZero introduces a Rapid Response GraphQL API for programmatic access to time-sensitive advisory and asset intelligence.

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

01 **Rapid Response GraphQL API** NEW 45

A new GraphQL API for Rapid Response lets users query advisories, list affected assets, annotate remediation status, and export data programmatically, giving programmatic access to time-sensitive attack surface intelligence alongside the existing Rapid Response Details UI (Asset Breakdown, Relevant Assets table, and Take Action workflow for running tests, adding notes, or creating tickets).

.
.
.
.

— Names GraphQL API and its operations but no endpoint or schema details.

snapshot-20260819

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ 1 RELEASE · 2026-08-19 CHANGELOG ↗

Semgrep is a static analysis tool that scans code for bugs, security vulnerabilities, and policy violations using customizable rules.

> Semgrep's latest release adds the ability to batch multiple findings from the same rule into a single Autofix pull or merge request, alongside an experimental v2 API for Agentic Workflow issues and several Rules & Policies UI refinements, while removing the standalone Code Autofix toggle from Settings.

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

01 **Rule deletion confirmation in Semgrep Editor** IMPROVED 65

Adds a confirmation dialog when deleting a rule in Semgrep Editor that shows which policies use the rule, automatically removes the rule from those policies, and deletes any policy left with no remaining conditions.

– Full before/after mechanism and named UI surface described changelog-20260819-746e8696

THINNER COVERAGE BELOW

02 **Experimental v2 API for Agentic Workflow issues** NEW 45

Exposes issues created by Semgrep Agentic Workflows through the experimental `v2 API`.

– Names the API version but no endpoint path or schema given changelog-20260819-746e8696

03 **Advisory affected-projects list** NEW 45

Adds an advisory affected-projects list directly below the description on the Rules & Policies > Advisories Details page.

– Exact UI navigation path given, but limited mechanism detail

changelog-20260819-746e8696

04 **Sort suggested memories by findings count** NEW 45

Enables sorting suggested memories by number of affected findings on Rules & Policies > Memories.

– Exact UI path given but thin on further mechanism changelog-20260819-746e8696

05 **Faster Dependencies page load** IMPROVED 20

Improves load time of the Dependencies page for organizations with large dependency sets.

— No numbers or mechanism given for the performance improvement

changelog-20260819-746e8696

└──▷ BREAKING ON UPGRADE

! The Code Autofix toggle has been removed from Settings. Supply Chain retains its own toggle for Upgrade Guidance and Supply Chain Autofix.

WAS THIS USEFUL?  

GitHub

Feedback

3 RELEASES · 2026-08-18 → 2026-08-19

BLOG ↗

GitHub is a web-based platform for version control and collaborative software development using Git repositories.

> GitHub shipped trend visibility for Code Quality findings, enterprise-managed controls for Copilot in JetBrains covering plugins, MCP servers, and OpenTelemetry, and token-type-specific credential revocation and deauthorization now available at the organization level.

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

01 **Enterprise managed settings for Copilot in JetBrains** NEW

90

Adds `enabledPlugins` to require Copilot plugins be enabled or disabled across JetBrains IDEs, `extraKnownMarketplaces` to make approved plugin sources available, and `strictKnownMarketplaces` to restrict plugin installation to approved sources only. Adds `allowedMcpServers` and `deniedMcpServers` to centrally control which MCP servers developers can connect to, and `permissions.disableBypassPermissionsMode` to prevent the Copilot agent from using Bypass Approvals or Autopilot. Also enables centrally managed OpenTelemetry configuration (collector endpoint, protocol, service name, resource attributes, content-capture policy) that takes precedence over developer settings, with the applied configuration exposed to developers under Settings > Tools > GitHub Copilot > Chat > OpenTelemetry.

Lock JetBrains Copilot users to approved MCP servers only, blocking any unlisted connections enterprise-wide.

```yaml
allowedMcpServers:
  - https://mcp.example.com/approved
deniedMcpServers:
  - https://mcp.untrusted.io
```

Prevent developers from enabling Bypass Approvals or Autopilot in the Copilot agent.

```yaml
permissions:
  disableBypassPermissionsMode: disable
```

— Every managed setting named verbatim with config examples and UI path

Enterprise managed settings in GitHub C…

02 **Trends tab in organization Code Quality dashboard** NEW

70

Adds a Trends tab to the organization-level Code Quality dashboard, plotting open findings over 7, 14, or 30-day windows grouped by health score or severity. Includes most improved and needs support repository ranking tables showing net change in open findings since the start of the selected period, and both the graph and tables respect repository filters applied at the top of the dashboard. Available on GitHub Enterprise Cloud, GitHub Team plans with GitHub Code Quality enabled, and GitHub Enterprise Cloud with data residency; not available on GitHub Enterprise Server.

— Names UI surface and windows but no API or exact steps

Track organization code quality trends …

03 **Token-type-specific credential revocation and deauthorization**

NEW

65

Adds token-type-specific bulk deauthorization via enterprise REST APIs and UI, letting enterprise owners revoke all SSO authorizations for a specific credential type (Personal access tokens, SSH keys, OAuth app tokens, or GitHub App user access tokens) across the enterprise or for a single user without touching other credential types. Adds matching bulk revocation to delete or revoke all user-level credentials of a specific type for an individual EMU user while leaving other types intact, and extends both actions from enterprise-only to the organization level via organization REST APIs and web UI for incident-response parity. All actions are captured in the audit log and trigger email notifications to affected users.

— Names token types and scope but no specific endpoint paths

Credential revocation and deauthorizati…

WAS THIS USEFUL? 👍 👎

noir

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

Hunt every Endpoint in your code, expose Shadow APIs, map the Attack Surface.

> Noir's latest releases focus on scan reliability and broader framework coverage: a new `--strict` mode and `errors` reporting surface previously silent analyzer failures, HTTP QUERY method (RFC 10008) support lands across ~24 frameworks, and a dozen new language/framework analyzers extend endpoint discovery.

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

01 Strict mode and scan-integrity error reporting NEW

95

Noir adds an `errors` key to JSON output that records per-tech analyzer failures and skipped files, distinguishing an empty result from an analyzer that never ran; skipped files are tallied per tech with up to five example paths instead of one entry each. The `errors` key propagates to `-f yaml` output and is mapped to `runs[0].invocations[0].executionSuccessful` in `-f sarif` output. A new `--strict` flag exits with code 2 after a degraded scan (analyzer failure or skipped files), writing the report before exiting, so failures and skipped files that were previously silently dropped now surface in CI.

Fail CI when any analyzer was skipped or errored, while still capturing the full JSON endpoint report for further processing.

```
$ noir scan . -f json --no-log --strict > endpoints.json
```

Fail CI when any analyzer errors or file is skipped — catches silent coverage gaps in regulated pipelines.

```
$ noir -b <source_dir> --strict
```

— Names exact keys, format mappings, exit code and runnable CI commands. v1.3.0

02 HTTP QUERY method (RFC 10008) support NEW

75

Noir adds end-to-end support for the HTTP QUERY method across roughly 24 frameworks including Express, Fastify, Hono, NestJS, Flask, aiohttp, Rails, Laravel, Phoenix/Plug, Echo, Fiber, Go `net/http`, axum, Ktor, http4k, Micronaut, Vert.x, JAX-RS/Quarkus/Dropwizard, ASP.NET Core, and Kemal, including OAS3 serialization and HTML report rendering.

— Lists every supported framework and output format, but no direct usage example. v1.3.0

THINNER COVERAGE BELOW

03 **New framework and language analyzers** NEW 50

Noir adds endpoint analysis support for twelve new frameworks: ServiceStack, Kratos, Helidon MP, Helidon SE, Feathers.js, Oak (Deno), Sails.js, Phalcon, CherryPy, Masonite, Padrino, and LoopBack 4.

— Names every new framework but gives no mechanism or usage detail. v1.3.0

04 **iOS deep-link route recovery** IMPROVED 15

Noir now recovers iOS deep-link routes that were previously missed during config-only scanning.

— Bare statement of a fix with no mechanism or example given. v1.3.0

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ 2 RELEASES · 2026-08-18 → 2026-08-19 BLOG ↗

Endor Labs is a software supply chain security platform that detects vulnerabilities and malicious dependencies in open source code.

> Endor Labs extended its AI SAST engine with full C language support and buildless scanning, launched an Agents Hub and Agent Governance control plane for AI coding agents, and open-sourced reproducible Endor Patch examples with Maven/Gradle install instructions.

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

01 **Open-sourced Endor Patches with build and install instructions**

NEW 90

Endor Labs open-sourced two complete Endor Patch examples in the endor-patches repository, including reproducible builds, attestations, and auditable diffs. Patched JARs in `patches/<library>/artifacts/` can be installed to a local Maven repository via `mvn install:install-file` and declared using the `-endor-YYYY-MM-DD` version string, with transitive override support via `<dependencyManagement>` in Maven or `resolutionStrategy` in Gradle so nothing else in the build needs to change.

Independently verify a patch's security diffs by inspecting the attestation file and matching upstream commit SHAs before installing.

```
$ cat patches/com.fasterxml.woodstox/woodstox-
core/security_attestation.json | jq '.[] | select(.fix_type ==
"FIX_TYPE_VULN")'
```

Reproduce the patched artifact from scratch and verify its checksum matches the published artifact.

```
$ # Extract the source tarball, then build with Docker and Bazel
per the included README.md
tar -xf patches/com.fasterxml.woodstox/woodstox-
core/source.tar.gz
cd woodstox-core-patched && docker run --rm -v "$PWD":/build
<hermetic-image> bazel build //...
sha256sum output/woodstox-core-6.3.1-endor-2024-09-03.jar
diff <(sha256sum output/woodstox-core-6.3.1-endor-2024-09-03.jar
| awk '{print $1}') SHA256SUM.txt
```

Install the patched JAR into your local Maven repository and override a transitive dependency so nothing else in the build needs to change.

```
$ mvn install:install-file \
  -Dfile=patches/com.fasterxml.woodstox/woodstox-
core/artifacts/woodstox-core-6.3.1-endor-2024-09-03.jar \
  -DgroupId=com.fasterxml.woodstox \
  -DartifactId=woodstox-core \
  -Dversion=6.3.1-endor-2024-09-03 \
  -Dpackaging=jar
```

— Exact file paths, commands, and version format make this directly runnable

When you can't upgrade: open source exa…

02 **C language support in AI SAST** NEW 75

Endor Labs AI SAST now supports C with buildless scans that run in the IDE and on pull requests, requiring no compiler hook or reproducible build. It pairs a deterministic call-graph and reachability engine with LLM reasoning to trace data flow across functions and files, catching buffer overflows, use-after-free, integer overflows, and other cross-function memory-safety flaws that pattern-based rules miss. C SAST runs alongside existing C SCA in the same platform, covering first-party C code and its open-source dependencies together, and is surfaced through AURI by Endor Labs as an independent security harness with policy enforced across any agent, model, or CI stage.

— Mechanism and scope well described; no exact CLI command or flag given

Why Endor Labs AI SAST for C finds what…

03 **Agent Governance control plane for coding agents** NEW 75

Agent Governance is a new inventory and control plane for AI coding agents, models, MCP servers, and skills running across developer workstations and cloud environments, with real-time monitoring of agent actions such as commands, file access, and tool calls. It adds policy enforcement that can block, alert on, or require approval for risky agent behavior before execution, backed by a built-in policy catalog extensible with custom rules, plus reusable, saveable filters for findings and projects under Coding Agent Governance.

— Named monitoring and policy mechanisms but no config keys or API given product docs

THINNER COVERAGE BELOW

04 **Agents Hub for AI agent discovery** NEW 55

A new Agents Hub catalog lets practitioners browse, filter by category, and select published Endor Labs security agents covering workflows such as finding triage,

dependency remediation, and scan-failure diagnosis. Teams can also monitor usage of these agents, including task counts and last-run time.

— UI catalog described with named workflows but no exact navigation path `product docs`

WAS THIS USEFUL? 👍 👎

GitLab

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

GitLab is a web-based DevOps platform providing Git repository hosting, CI/CD pipelines, and collaboration tools for software development teams.

> GitLab shipped a Git Clone Override Policy that lets repos declare automatic shallow/partial clone behavior via a policy file and companion binary, alongside expanded usage visibility for personal access tokens.

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

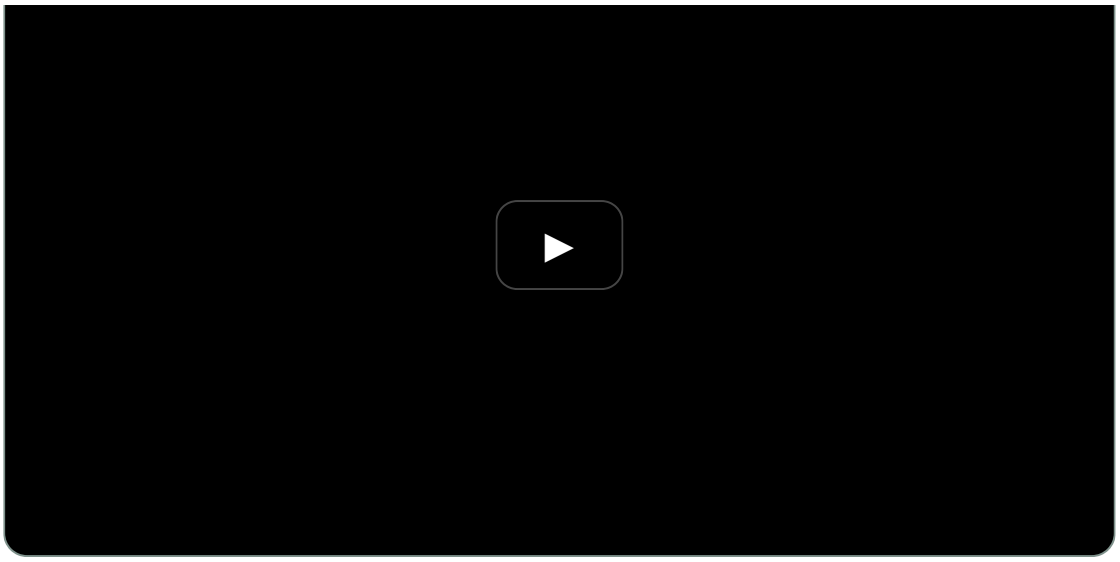
01 Git Clone Override Policy for automatic optimized clones NEW

98

A new self-contained Go binary `git-clone-override-policy` intercepts bare `git clone <URL>` calls and runs a fixed 13-step shallow-and-partial-clone sequence based on a `.afullhistorycloneoverridepolicy.toml` (`schema_version = 2`) file stored in the repo root, before enforcing any policy on disk. The policy file supports `[git_config]` (scoping `core.compression` , `http.postBuffer` , `http.lowSpeedLimit` , `http.lowSpeedTime` , `pack.windowMemory` , `pack.packSizeLimit` , `pack.threads` to the cloned repo only), `[fetch]` flags such as `--depth=1` and `--filter=blob:none` , `[sparse_checkout]` with `conemode = 'auto'` (resolved to `cone` / `no-cone`), `includes` , `excludes` , and `exclude_extensions` covering 30+ binary file types, plus a `[policyinfo]` section with `description` and `OptimalForAIAgents` metadata. The binary runs on Linux, macOS and Windows (amd64/arm64, including CMD, PowerShell, Git Bash, WSL) with no dependency beyond `git` on PATH, and is bypass-safe: clones that already pass a flag like `--depth` or `--filter` , or repos without a policy file, pass through unchanged.







Commit a policy file to a monorepo so every bare clone — by developers, CI jobs, or AI agents — automatically gets a shallow partial clone that skips binaries, without anyone remembering the flags.

```toml
# .afullhistorycloneoverridepolicy.toml at repo root
schema_version = 2

[policyinfo]
description = "Optimized clone for CI and AI agents"
OptimalForAIAgents = true

[git_config]
"core.compression" = 0
"http.postBuffer" = "1024M"
"pack.threads" = 4

[fetch]
flags = ["--depth=1", "--filter=blob:none"]

[sparse_checkout]
enabled = true
conemode = "auto"
includes = ["/*"]
exclude_extensions = [
  "png", "jpg", "gif", "svg", "pdf",
  "zip", "tar", "gz", "mp4", "exe",
  "dll", "woff", "woff2", "ttf", "psd"
]
case_variants = "both"
```

Run the override binary in place of `git clone` so the policy is enforced for a one-off zero-footprint clone without any global git config changes.

```
$ ./git-clone-override-policy https://example.com/group/my-repo.git
```

— Full config schema, binary mechanism, and runnable examples all provided.

Avoid the massive end-to-end tax of def…

THINNER COVERAGE BELOW

02 **Personal access token usage tracking in token detail view**

IMPROVED

50

The Personal Access Token detail view now surfaces last-five connection IPs and per-scope usage timestamps alongside existing scopes, assigned roles, and expiration dates; usage times refresh every 10 minutes via the API.

— UI location and refresh interval given, but no endpoint or field name. product docs

WAS THIS USEFUL? 👍 👎

Black Duck

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

Black Duck is a software composition analysis tool that identifies and manages open source components, vulnerabilities, and licensing risks in applications.

> Black Duck's Polaris platform added AI-assisted SAST fix pull requests, two-way bug tracker sync across five systems, secure tunnels into self-hosted SCMs and Jira, organization-wide rule governance, and broadened static and SCA engine coverage.

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

01 AI-assisted SAST fix pull requests NEW 82

Polaris can now generate AI-assisted fix pull requests for SAST findings both CI-natively via the Bridge CLI — triggerable from GitHub Actions, GitLab CI, Bitbucket Pipelines, Jenkins, or Azure DevOps pipeline configuration with no server-side scheduling — and on-demand or policy-driven from the UI, individually or in bulk, or via issue policy actions. Admins can cap open PRs per branch or per run and filter by `CWE`, issue type, language, and file pattern; a new View issue policies action shows which policies, compliance requirements, and governance rules are attached to a given SAST finding.

— Names triggers, CI systems and filters but no exact command

Polaris release update: Two-way bug tra…

02 Central rule governance for scan policies NEW 82

Security teams can define rule profiles at organization, application, project, or branch level with hierarchical inheritance, enforced automatically across CI pipelines, CLI scans, and integrations — for example, setting a mandatory organization-level profile via Settings > Rule Profiles so every scan picks it up without per-pipeline changes. The feature adds RBAC, audit trails, and version awareness, with visibility into which profile was used for each scan.

Apply a mandatory central rule profile at the organization level so every scan — CI, CLI, or integration — enforces the same ruleset without per-repo YAML.

📍 # In Polaris UI: Settings › Rule Profiles › New Profile # Set scope: Organization # Enable Mandatory and configure inheritance for application/project/branch overrides # All pipelines pick up the profile automatically on next scan — no per-pipeline changes needed.

— Includes exact UI navigation and inheritance scope

Polaris release update: Two-way bug tra…

03 Two-way bug tracker sync across five systems NEW 81

Polaris now syncs triage states, ticket statuses, and fix-by dates bidirectionally with GitLab Issues, GitHub Issues, Jira, ServiceNow, and Azure Boards — changes in either system propagate automatically, including fix-by date updates flowing to Jira/ServiceNow/Azure Boards due dates and back. GitLab Issues, GitHub Issues, and ServiceNow are new native bug-tracking targets supporting individual and bulk export, single-ticket bundling for ServiceNow, and policy-driven issue creation; Jira gains configurable field mapping to standard or custom fields and customizable ticket descriptions. Per-project bug-tracker integration mappings can be saved as reusable templates across projects and workflows, scoped to all branches or the default branch only, and linked tickets now auto-close when findings are dismissed or absent in a later scan, with a reappearing issue creating a new ticket to preserve the audit trail.

— Rich detail on five systems and behavior, no config surface named

Polaris release update: Two-way bug tra…

---

04 **Rapid Scan Static engine expanded checks** NEW 73

The Rapid Scan Static engine extends its seven existing Java JWT checks (algorithm confusion, nonexpiring tokens, ignored expiration/start times, missing claims validation, untrusted decoding) to Scala, and adds six new Scala web security checks covering permissive and credentialed CORS origins, excessive preflight age, missing TLS, and disabled security headers across the http4s and Play frameworks. It also adds a crypto bits private key pattern to detect embedded cryptographic private keys in source and config files, and adds support for `Go 1.26` and `Dart 3.11`.

— Names exact checks and language versions, no CLI flag given

Polaris release update: Two-way bug tra…

---

05 **fAST SCA package manager coverage expanded** IMPROVED 73

The fAST SCA engine adds new detectors for UV (Python), Rush monorepo tooling, a tree-based Conda detector, and Ant/Ivy dependency extraction. Ecosystem coverage is extended to Poetry, pnpm, SBT, SetupTools, Bazel 9, `.slnx` solution files for NuGet, buildless `go.mod` analysis for Go, and Cargo workspace support with include and exclude controls.

— Names every package manager and file format added

Polaris release update: Two-way bug tra…

---

06 **Secure tunnels to self-hosted SCMs and Jira** NEW 68

Polaris adds outbound-only secure tunnel connectivity — requiring no inbound firewall changes — to self-hosted Azure DevOps Server, GitLab Self-Managed, Bitbucket Server and Data Center, and on-premises Jira, built on the Bridge CLI and Teleport infrastructure. Multitunnel support for on-prem Jira connectivity is restricted to org admins via RBAC, with full audit logging.

— Names target systems and stack, no setup steps given

Polaris release update: Two-way bug tra…

THINNER COVERAGE BELOW

07 **Webhook-triggered scanning on self-hosted SCMs** NEW

57

SAST and SCA scans can now be triggered by webhooks on push, branch creation, and pull/merge request activity for self-hosted SCMs, with PR/MR decoration that includes severity breakdowns and Code Insights reports for Bitbucket.

— Describes trigger events and Bitbucket reporting, no config keys

Polaris release update: Two-way bug tra…

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ 1 RELEASE · 2026-08-19 CHANGELOG ↗

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

> Veracode extended eLearning reporting to support longer historical date ranges.

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

01 **Extended date range for eLearning reports** IMPROVED

35

eLearning reports now support a start date up to one year in the past, enabling longer historical reporting windows.

— Names the limit but no mechanism or UI path changelog-20260819-1ab01652

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ 1 RELEASE · 2026-08-19 DOCS ↗ BLOG ↗

Snyk is a developer security platform that finds and fixes vulnerabilities in open source dependencies and container images.

> Snyk's Evo platform replaces its Risk Index with a 0–1,000 model risk score and adds Continuous Offensive Security, alongside new compliance framework mappings, custom attack-category policies, a `snyk fix` CLI command, and a Remediation Agent.

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

01 Model risk score replaces Risk Index BREAKING

85

Evo by Snyk replaces the previous Risk Index with a new model risk score ranging 0–1,000 that combines attack success rate with impact weighting, tested across prompt injection, data exfiltration, and insecure code generation attack categories. The score is split into four severity bands for at-a-glance triage: low (0–249), medium (250–499), high (500–749), and critical (750–1,000). Policies built on the old Risk Index remain visible but no longer evaluate under the new model risk score — existing Risk Index-based policies must be reviewed and recreated against it to keep enforcement active.

– Names ranges, bands and migration note but no exact recreation steps

2026-08-19-1e4a5632

02 Continuous Offensive Security in Evo NEW

60

Introduces Continuous Offensive Security (COS), a new offensive scanning capability covering Targets, Scans, and Findings and Reports, with dedicated documentation for defining target scope, configuring authentication, managing targets, monitoring scans, and interpreting results.

– Describes scope and doc areas but no config keys or commands 2026-08-19-1e4a5632

03 `snyk fix` CLI command NEW

60

Adds a `snyk fix` CLI command, documented at `commands/fix.md`.

– Names exact command and doc path but no usage behaviour 2026-08-19-1e4a5632

THINNER COVERAGE BELOW

04 Compliance framework mappings for model risk findings NEW

55

Maps model risk findings to OWASP LLM Top 10, OWASP Agentic Security, MITRE ATLAS, and NIST AI 600-1/100-2e2025 compliance frameworks via new Framework mappings documentation.

— Names frameworks but not how mappings are surfaced or used 2026-08-19-1e4a5632

05 **Custom policies for attack categories and goals in Evo** NEW

45

Enables custom policies in Evo scoped to specific attack categories or goals, in addition to automatically applied default policies across common attack categories.

— Explains scoping concept but gives no config syntax or example 2026-08-19-1e4a5632

06 **Remediation Agent capability** NEW

45

Introduces a Remediation Agent capability, documented at `agent-security/remediation-agent/remediation-agent.md`.

— Names doc path only; behaviour not described 2026-08-19-1e4a5632

└──▷ **BREAKING ON UPGRADE**

! Policies built on the old Risk Index remain visible but no longer evaluate under the new model risk score — existing Risk Index-based policies must be reviewed and recreated against the new model risk score to keep enforcement active.

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

Cloudsmith is a cloud-based software supply chain platform for securely managing and distributing software packages and artifacts.

> Cloudsmith expanded Go upstream configuration to support any GOPROXY-compatible module proxy, moving beyond the fixed public mirror.

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

01 **GOPROXY-compatible Go upstream sources** IMPROVED

50

Go upstream sources in Cloudsmith can now be configured to point to any GOPROXY-compatible module proxy, expanding beyond the previously supported public mirror at `proxy.golang.org`.

— Names the config surface but no steps or UI path

Configure Go upstreams for any GOPROXY-…

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

DataDome

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

DataDome is a bot management and API protection platform that defends applications against automated attacks and fraud using AI-driven detection.

> DataDome's Magento module gained a new configuration property to enable MCP support.

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

01 **MCP support flag in Magento module** NEW 50

The Magento module adds a `datadome_enable_mcp_support` property to its configuration to enable MCP support.

– Names exact config key but no detail on MCP behavior.

Magento module 1.2.247 (2026-08-19)

WAS THIS USEFUL? 👍 👎

Cure53 DOMPurify

Feedback ⓘ 1 RELEASE · 2026-08-19 NOTES ↗

DOMPurify - a DOM-only, super-fast, uber-tolerant XSS sanitizer for HTML, MathML and SVG. DOMPurify works with a secure default, but offers a lot of configurability and hooks. Demo:

> DOMPurify 3.4.14 expands its default SVG allow-list to cover two more presentation attributes.

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

01 **SVG pointer-events and vector-effect allow-listed** IMPROVED 60

DOMPurify's built-in SVG allow-list now includes the `pointer-events` and `vector-effect` presentation attributes, letting sanitized SVG content use them without custom configuration.

– Names exact attributes added but no mechanism detail 3.4.14

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-08-18 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 gains reserved-machine API selection, platform-aware demux routing, and Proxmox read-retry logic

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-18
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-18
```

Validates explicitly named machines against db.list_machines(include_reserved=True) so callers can submit file, URL, and download-and-execute tasks to a reserved machine by name without the API rejecting it.

Defers platform assignment for archive/container submissions until after demultiplexing, then routes each extracted file's tasks to machines compatible with the detected platform; `machine='all'` expansion now uses the extracted platform.

Adds `_get_with_retry` helper to the Proxmox machinery module that retries `proxmoxer.ResourceException` and `requests.RequestException` on VM enumeration, snapshot enumeration, task-status polling, and post-task VM-status reads, bounded by the existing machinery timeout.

Aligns BSON and Protobuf netlog argument parsing by name and converter rather than positional index, with isinstance(int) simplification for the integer converter while preserving 64-bit signed-int semantics.

Adds colored, structured log output (`ColoredFormatter`) to `utils/gcp_pubsub_service.py` with pre-compiled regex patterns highlighting worker names, task IDs, and analysis IDs.

▾ 1 more release in this issue · 2026-08-18 → 2026-08-19

changes since 2026-08-17 CODE ↗

CAPEv2 adds a `task_id_1` MongoDB index on the `calls` collection and ships a Tracee Docker setup script for Linux agent deployments.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-17
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-17
```

└──▷ **TRY IT**

Bootstrap a new Linux analysis guest with all agent dependencies (including Tracee and pyinotify) in one step.

```
$ sudo bash extra/linux_agent.sh && sudo bash extra/tracee.sh
```

Adds `task_id_1` index on the MongoDB `calls` collection (field `task_id`, ascending) during standalone processor startup, eliminating full collection scans on task-scoped delete and replace operations.

Adds `exit_on_connection_failure` parameter to check_webgui_mongo(), allowing the standalone processor (`utils/process.py`) to log a warning and continue when MongoDB is temporarily unavailable instead of hard-exiting.

Reconciles required MongoDB indexes at standalone processor startup by calling check_webgui_mongo(exit_on_connection_failure=False) early in main(), so existing installations receive missing indexes automatically on upgrade.

Adds `extra/tracee.sh` — a new setup script that installs Docker CE and pulls `aquasec/tracee:0.24.1` for Linux agent deployments.

Updates `extra/linux_agent.sh` to install `python3-pyinotify`, `systemtap-runtime`, `ca-certificates`, `gnupg`, and `lsb-release` via `apt` instead of `pip3 install pyinotify`, consolidating Linux agent dependencies into a single package-managed step.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

ALEAPP

Feedback ⓘ 1 RELEASE · 2026-08-18 NOTES ↗ CODE ↗

Android Logs Events And Protobuf Parser

> ALEAPP added a detailed Proton Mail MailX parser with four artifact types, five new app parsers (Instagram, Tinder, Netflix, Reddit, Amazon Shopping), and a wide set of enhancements to WhatsApp, Snapchat, TikTok, Wire, Chrome, Samsung My Files, GroupMe, and LINE artifact extraction.

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

01 **Proton Mail MailX Room store parser** NEW 90

New `protonmailDbMail.py` parser surfaces four artifacts from Proton Mail's db-mail (MailX) Room store: MailX Messages (time, subject, sender/recipient addresses decoded from JSON, read/replied/forwarded flags, attachment count, size, conversation ID, message ID), MailX Attachments (name, size, MIME type, disposition, saved content URI from `MessageAttachmentMetadataEntity`), MailX Contacts (`ContactEmailEntity` rows), and MailX Account (`UserEntity` + `AccountEntity` including `createdAtUtc` in Unix millis).

— Names parser file, entities, and exact fields extracted v2026.3.1

02 **TikTok account and activity artifacts** NEW 60

Adds a TikTok Account artifact and TikTok activity artifacts, with multi-account database handling for IM data and deduplication of duplicate extraction paths.

— Describes multi-account and dedup mechanism, not just app name v2026.3.1

THINNER COVERAGE BELOW

03 **Samsung My Files artifacts** NEW 49

Adds seven new Samsung My Files artifacts pulled from previously unparsed databases, with cached thumbnails now displayed inline in local file listings.

— Gives count and describes new inline thumbnail behavior v2026.3.1

04 **GroupMe artifacts** NEW 46

Adds a GroupMe FCM queued messages artifact, message deletion state reporting, and scoped group-creator lookup.

— Names three distinct GroupMe additions without deeper mechanism v2026.3.1

05 **WhatsApp add-on artifacts** NEW 45

Adds six new WhatsApp add-on artifacts: reactions, edits, revokes, polls, receipts, and system events.

– Lists six named artifact types but no extraction mechanism v2026.3.1

06 **Wire data recovery for encrypted app generations** NEW

42

Adds Wire data recovery, recovering data that survives across encrypted app generations.

– States mechanism (surviving app generations) but no further specifics v2026.3.1

07 **Chrome payment artifacts** NEW

37

Adds a Chrome saved credit cards artifact and a Chrome payments customer data artifact.

– Names two artifact types with no further detail v2026.3.1

08 **LINE message and call reporting without contacts** IMPROVED

36

Adds LINE message and call reporting that works even when contact tables are empty.

– Describes edge-case handling but no further mechanism v2026.3.1

09 **New app parsers: Instagram, Tinder, Netflix, Reddit, Amazon Shopping**

NEW

35

Adds Android parser support for Instagram, Tinder, Netflix, Reddit, and Amazon Shopping.

– Five named apps but each described only as bare support added v2026.3.1

10 **Snapchat chat media and session artifacts** NEW

33

Adds Snapchat chat media, `core.db`, and session artifacts.

– Names a specific database file but little other detail v2026.3.1

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ 1 RELEASE · 2026-08-18 NOTES ↗ CODE ↗

iOS Logs, Events, And Plist Parser

> iLEAPP added a new logarchive dialed-number call artifact with detailed CommCenter/ContactSearchManager parsing, rolled out new parsers across nine apps (Snapchat, Tinder, TikTok, Reddit, OneDrive, Google Drive, WhatsApp, Chrome, GroupMe), reworked Proton Mail's inbox into a conversation view, and made Chromium database parsing more resilient.

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

01 **Proton Mail inbox conversation view** IMPROVED 85

Upgrades Proton Mail Inbox Messages parsing to a conversation view with inline media, adding `remote_conversation_id`, Conversation Subject, From Me direction, and first cached Attachment columns via a new `data_views.conversation` block.

— Names the data block and every new column added v2026.3.1

02 **Logarchive dialed-number call artifact** NEW 70

Adds a new `logarchive dialed numbers` artifact that extracts CommCenter `call.provider` setup/teardown blocks, Call(StatusUpdate) state chains, and `ContactSearchManager` digit-sequence entries from unified log archives.

— Names exact log structures parsed but no run command shown v2026.3.1

03 **New app artifact parsers across nine apps** NEW 60

Adds parsers for new apps and data types: Snapchat (including a Memories Search Index), Tinder, four TikTok artifacts (Watch History, Account, Published Videos, App Sessions), Reddit (user and subscription data), OneDrive, Google Drive DriveKit artifacts, four WhatsApp add-on artifacts (reactions, channel tallies, group events, membership changes), Chrome saved credit cards and Chrome payments customer data artifacts, and GroupMe message deletion state reporting.

— Lists many named artifacts but no extraction mechanism given v2026.3.1

THINNER COVERAGE BELOW

04 **Chromium database error handling** IMPROVED 35

Artifacts now skip Chromium databases that cannot be opened instead of aborting the entire artifact, improving resiliency when parsing corrupted or locked database files.

— Describes behavior change but no config or flag involved v2026.3.1

WAS THIS USEFUL?

NETWORK SECURITY

◆ Firewalls & Network Defense

Aviatrix CoPilot

Feedback ⓘ 1 RELEASE · 2026-08-19

CHANGELOG ↗

Aviatrix CoPilot is an AI-powered assistant that helps operators manage, troubleshoot, and optimize cloud network infrastructure across multiple cloud providers.

> Aviatrix CoPilot 4.36.0 adds new visibility columns to the gateway security rules tab.

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

01 **Session Count and Logging columns in security rules tab** NEW

48

CoPilot's security rules tab for gateways now includes Session Count and Logging columns to improve traffic visibility.

— Names exact UI columns but no further mechanism given changelog-20260819-01a55e4e

WAS THIS USEFUL?

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Wiz

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

Wiz is a cloud security platform that identifies and remediates risks across cloud infrastructure and applications.

> Wiz brought Penetration Test Findings to General Availability, consolidating offensive testing results into one platform tied to real-time cloud context and exposure management workflows.

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

01 **Penetration Test Findings GA with CTEM workflow integration**

NEW 50

Wiz Penetration Test Findings is now Generally Available, giving a single home for pen-test results from bug bounties, third-party audits, internal red team exercises, and AI pen-test scanning, connected to real-time cloud context on the Wiz Security Graph. Findings are automatically linked to continuous prioritization, ownership assignment, and remediation tracking, and the feature bridges the Validation stage back to Discovery, Prioritization, and Mobilization to support Continuous Threat Exposure Management (CTEM) workflows.

— Describes scope and workflow tie-ins but no concrete UI paths or config

2026-08-19-d958d675

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Wiz Cloud

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

Wiz Cloud is a cloud security platform that identifies and remediates risks across cloud infrastructure, applications, and data.

> Wiz brought Pen Test Findings to general availability with new Security Graph correlation, integrated Google Threat Intelligence across risk, malware, and AI-agent findings, and widened Security Graph coverage into hybrid/on-prem, third-party inventory, and SaaS environments.

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

01 **Google Threat Intelligence integration across finding types**

NEW 70

Wiz integrates Google Threat Intelligence, enriching risk and threat findings with Mandiant frontline intelligence, VirusTotal crowdsourced intelligence, and Google threat insights, available out of the box to all Wiz customers. The enrichment extends to malware findings for deeper investigation context beyond a simple verdict, and to malicious AI Skill intelligence that identifies potentially risky open-source skills AI agents can import, addressing agentic AI attack surfaces.

— Names all three enrichment sources and finding types, but no config or endpoint to act on

2026-08-19-8dbcb4c7

02 **Pen Test Findings GA with Graph Controls** NEW 60

Pen Test Findings reaches general availability with new Graph Controls that surface attack paths tied to pen-test findings, enabling blast-radius visualization correlated against external network exposure, identity privileges, secrets, data, and lateral movement paths within the Wiz Security Graph.

— Explains mechanism and correlated risk types but no UI/API path given

2026-08-19-8dbcb4c7

THINNER COVERAGE BELOW

03 **Security Graph extended to hybrid and on-prem via Sensor Workload Scanner**

NEW 50

The Wiz Security Graph now extends into on-premises and hybrid environments via the Sensor Workload Scanner, giving a unified hybrid security posture view.

— Names the scanner mechanism but not setup steps 2026-08-19-8dbcb4c7

04 **Third-party inventory and findings connected into Security Graph**

NEW

45

Wiz now connects third-party inventory and risk findings into the Security Graph, enabling unified vulnerability prioritization across all finding types using Wiz context and attack paths.

— States capability and benefit but no specific integrations named 2026-08-19-8dbcb4c7

05 **Security Graph coverage extended to SaaS applications**

IMPROVED

40

Security Graph risk assessment coverage now extends to SaaS applications including M365, closing visibility gaps across enterprise SaaS.

— Names M365 but coverage scope otherwise unspecified 2026-08-19-8dbcb4c7

WAS THIS USEFUL?  

◆ Runtime Security

Docker Sandboxes

Feedback ⓘ 1 RELEASE · 2026-08-19 NOTES ↗

Docker Sandboxes provides isolated, lightweight environments for securely running and testing AI models and applications with controlled dependencies.

> Docker Sandboxes shipped declarative `.sbxenv.yaml` environments for reproducible sandbox provisioning, dynamic secret sources (1Password, AWS Secrets Manager, host commands) for sandbox credentials, and experimental NVIDIA GPU and USB device passthrough, alongside a new standalone `sbx` CLI, bulk sandbox pruning, and expanded environment-variable and kit-signing support.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 Declarative sandbox environments via .sbxenv.yaml NEW

100

Introduces `.sbxenv.yaml` (fallback `.sbxenv.yml`) files declaring a reproducible sandbox with `schemaVersion`, `name`, `agent`, `kits`, `workspace`, `additionalWorkspaces`, `env`, `sandboxOptions`, `secrets`, `bindings`, `registries`, `mcp`, and `ports` fields, managed via `sbx env run [PATH...]`, `sbx env create [PATH...]`, `sbx env exec [PATH...] -- COMMAND [ARG...]`, and `sbx env rm [PATH...]` (requires `sbx` 0.39.0+). Multiple files can be layered — `sbx env run base.sbxenv.yaml local.sbxenv.yaml` — merging nested mappings by key, concatenating lists, and letting later scalars override earlier ones for team-base-plus-local-override setups. `additionalWorkspaces` entries take `path` and `readOnly` to mount extra repositories; `sandboxOptions` supports `template`, `memory` (e.g. `8g`), `cpus`, `pullPolicy` (`always` / `missing` / `never`), and `profile`; `secrets` maps service names to `value`, `ref` (e.g. `op://Vault/Item/field`), or `command` sources with optional `refresh`, `backend` (`sdk` / `cli`), and `noVerify`; `bindings` approves credential-injection domains into the host-global `credentials.yaml` (removable with `sbx env rm --prune-bindings`); `registries` maps hostnames like `ghcr.io` to pull credentials; `mcp.servers` registers remote (`url`) or local stdio (`command` / `args`) MCP servers with the built-in gateway, persisting across `sbx env rm`; `ports` declares sandbox-to-host mappings with `sandbox`, `host`, `protocol`, and `hostIP`; and `workspace` takes an object form with `path` and `clone` (default `false`), overridable per-invocation with `--clone` / `--clone=false`.

Bootstrap a shared web-app sandbox with Playwright tooling and a published dev port — teammates run the same command to get an identical environment.

```yaml
# web-app-env/.sbxenv.yaml
schemaVersion: "1"
name: web-app
agent: claude
workspace: ./web-app
kits:
  - docker.io/sbx/playwright-kit:latest
env:
  NODE_ENV: test
ports:
  - sandbox: 3000
    host: 3000

# Then, from web-app-env/:
sbx env run
```

Layer a personal memory override on top of a committed team config without editing the shared file — useful when your machine has more RAM than the team baseline.

```
$ sbx env run base.sbxenv.yaml local.sbxenv.yaml
```

Run automated tests inside a sandbox in CI: create it headlessly, execute the test suite, then tear it down — secrets resolve on the runner, never stored in the file.

```
$ sbx env create
sbx env exec -- npm test
sbx env rm --force
```

Commit a `.sbxenv.yaml` with your project so any contributor can reproduce the exact same sandbox environment with a single command.

```
$ sbx env run
```

— Every config field and command enumerated with runnable examples v0.39.0

02 Dynamic secret sources for sandbox credentials NEW

96

Adds `--ref` to register a dynamic secret source using a 1Password reference (`op://...`) or AWS Secrets Manager ARN (`arn:aws:secretsmanager:...`), resolved at runtime by the host CLI rather than stored as plaintext, and `--command` to

register any host tool that prints a secret to stdout (e.g. `gh auth token`) as a dynamic source. `--refresh <duration>` controls caching (default 55 minutes), with `--refresh on-demand` resolving on every use. `--no-verify` stores a source that can't be resolved at registration time, and `--show-error` surfaces provider errors when verification fails (the two cannot be combined); `--ref` and `--command` are mutually exclusive and cannot combine with `--token`, `--oauth`, or `--registry`. `OP_ACCOUNT` and `AWS_PROFILE` environment variables select a specific 1Password account or AWS profile. On Linux hosts without a running Secret Service, credentials fall back to a file under the user config directory with restricted permissions and a warning that the secret is protected only by file permissions. `sbx secret set` and the experimental `sbx secret set-custom` gained the same dynamic-source, refresh, verification and error-output options for injecting credentials for known hosts via the host-side proxy without exposing values to the agent; third-party kits extending a built-in agent inherit credentials but not built-in provenance and require explicit approval.

Pull an Anthropic API key from 1Password at runtime so the real credential is never stored in the sandbox — useful for teams that rotate secrets centrally.

```
$ docker sandbox credentials set anthropic --ref
'op://Work/Anthropic/credential' --refresh 30m
```

Use the authenticated GitHub CLI on the host to supply a refreshing token for agents that open PRs or call GitHub APIs, avoiding any stored secret.

```
$ docker sandbox credentials set github --command 'gh auth token'
--refresh on-demand
```

Register an AWS Secrets Manager secret for OpenAI using a specific AWS profile, storing the ARN reference rather than the key value.

```
$ AWS_PROFILE=prod docker sandbox credentials set openai --ref
'arn:aws:secretsmanager:us-west-2:123456789012:secret:openai-api-
key'
```

Inject a GitHub token from the host CLI on every sandbox credential use without ever copying the token into a secret store.

```
$ github --command 'gh auth token' --refresh on-demand
```

Source Anthropic and GitHub credentials from 1Password on the host, keeping real values out of the sandbox.

```
$ github --ref 'op://Work/GitHub/token' anthropic --ref
'op://Work/Anthropic/credential'
```

Pull an Anthropic API key from AWS Secrets Manager using its ARN so the sandbox proxy resolves it via the authenticated AWS CLI on the host.

```
$ anthropic --ref 'arn:aws:secretsmanager:us-west-
2:123456789012:secret:anthropic-api-key'
```

— Full flag set, defaults and fallback behavior with runnable examples v0.39.0

---

03 **Experimental NVIDIA GPU passthrough via VFIO** NEW 96

Adds `--gpu` to `sbx create` and `sbx run` for NVIDIA GPU passthrough into sandboxes on x86_64 Linux hosts using VFIO, gated behind `sbx settings set platform.allowExperimentalFeatures true` and `sbx settings set feature.sandbox-gpu true`. A `prepare-nvidia-bundle.sh` script assembles the required EROFS driver bundle (distributed as `nerdbox-nvidia-modules-x86_64.zip`) and installs it at `/usr/libexec/nerdbox-nvidia-bundle.erofs`, with `OUTPUT` overriding the install path and `ACCEPT_NVIDIA_LICENSE=1` for non-interactive builds; hosts need IOMMU enabled in BIOS and `iommufd` / `vfio_pci` kernel modules loaded via `sudo modprobe -a iommufd vfio_pci`. After upgrading, the kernel modules must be rebuilt to match the new guest kernel by re-running `prepare-nvidia-bundle.sh` with the new release's module zip.

Enable the experimental GPU feature flag, then launch a sandbox with NVIDIA GPU passthrough for running GPU-accelerated AI workloads.

```
$ sbx settings set platform.allowExperimentalFeatures true
sbx settings set feature.sandbox-gpu true
sbx create --gpu claude .
```

— Exact flags, script, paths, and upgrade steps given product docs

---

04 **sbx CLI installation across macOS, Windows, and Linux** NEW 88

Introduces the `sbx` CLI for running AI coding agents in isolated microVMs without requiring Docker Desktop or Docker Engine. Installs on macOS via `brew trust docker/tap` then `brew install docker/tap/sbx`; on Windows via `winget install -h Docker.sbx`; on Linux alongside Docker Engine via `curl -fsSL https://get.docker.com | sudo SBX=1 sh`; on Linux without Docker Engine via

Docker's `apt` repository and `sudo apt-get install docker-sbx` (using `REPO_ONLY=1` with the convenience script); or manually from the `sbx-releases` repository binaries. Authentication is via `sbx login`, which opens a browser for Docker OAuth.

Install `sbx` on Linux alongside Docker Engine in one step — useful for a fresh CI runner or developer VM.

```
$ curl -fsSL https://get.docker.com | sudo SBX=1 sh
```

Install `sbx` on Linux without Docker Engine when you only need the sandbox runtime, then authenticate.

```
$ curl -fsSL https://get.docker.com | sudo REPO_ONLY=1 sh
sudo apt-get install docker-sbx
sbx login
```

Install `sbx` on macOS via Homebrew and sign in to start running AI agent sandboxes.

```
$ brew trust docker/tap && brew install docker/tap/sbx && sbx
login
```

— Every install path and command given, runnable end to end `product docs`

---

05 **Environment variable injection and persistence** NEW

86

Adds `-e` / `--env` and `--env-file` flags to `sbx run` and `sbx create` to set individual environment variables or load them from files (following `docker run` precedence rules), including copying a value from the host by naming a variable without a value (requires v0.39.0+); `sbx exec -e` and `sbx exec --env-file` set variables for a single exec session without altering the sandbox's stored environment. Variables can be persisted across future sessions by appending exports to `/etc/sandbox-persistent.sh`, which is sourced when Bash starts inside the sandbox.

Inject a debug log level and a host-side API URL into a sandbox at launch without baking them into the image.

```
$ export API_URL=https://api.example.com
sbx run -e LOG_LEVEL=debug -e API_URL --env-file .env.sandbox
claude
```

Persist a new variable across all future sessions of an existing sandbox without stopping it.

```
$ sbx exec -d my-project bash -c "echo 'export INTERNAL_API_URL=https://api.example.com' >> /etc/sandbox-persistent.sh"
```

Inject secrets and config into a sandbox at creation time, mirroring familiar `docker run` env-var precedence.

```
$ sbx run --env-file .env -e MY_TOKEN=abc123 my-image
```

— All flags, precedence rule and persistence file named with examples v0.39.0

06 **Experimental USB device passthrough** NEW 81

Adds an experimental `--usb` flag to `sbx create` (enabled via the `DOCKER_SANDBOXES_FEATURE_SANDBOX_USB` environment variable) to re-attach specified USB devices inside a sandbox via usbfs; supported on Linux x86_64/ARM64 only.

Pass through a specific USB device to a sandbox on Linux for hardware-in-the-loop agent testing.

```
$ DOCKER_SANDBOXES_FEATURE_SANDBOX_USB=1 sbx create --usb /dev/bus/usb/001/003 my-sandbox
```

— Names flag, env var, mechanism and platform scope with example v0.39.0

07 **Bulk cleanup of stopped sandboxes** NEW 75

Adds `sbx prune` to remove all stopped local sandboxes, with `--dry-run` to preview removals and `--filter since=<duration>` (e.g. `since=168h`) to exclude recently stopped sandboxes; a running sandbox is never removed.

Preview which stopped sandboxes would be pruned, excluding any stopped within the last week, before committing to removal.

```
$ sbx prune --dry-run --filter since=168h
```

— Flags, filter syntax and safety behavior with example v0.39.0

08 **Ollama provider selection for sandbox agents**

BREAKING

75

Adds `--provider ollama` flag to `sbx run` (e.g. `sbx run --model gemma4 --provider ollama claude`) to use an existing Ollama installation instead of the default provider. This flag replaces the former `ollama/` model-name prefix, which `sbx run --model` no longer accepts.

Run a Claude Code sandbox against a local Ollama installation using a specific model instead of the default cloud provider.

```
$ sbx run --model gemma4 --provider ollama claude
```

— Names flag, before/after syntax, and runnable example v0.39.0

09 **Kit signing and verification via Sigstore** NEW

62

Adds `sbx kit sign` and `sbx kit verify` to sign and verify kits with cosign-compatible Sigstore signatures, with optional policy enforcement applied at load time.

— Names commands and signing mechanism, no worked example v0.39.0

THINNER COVERAGE BELOW

10 **SANDBOX_NAME and SANDBOX_ID environment variables** DEPRECATED

57

Exposes `SANDBOX_NAME` and `SANDBOX_ID` environment variables inside sandboxes, matching the values shown by `sbx ls --json`; the `SANDBOX_VM_ID` environment variable is deprecated in favor of these two.

— Names variables and deprecation but no migration example v0.39.0

11 **Kit source allowlist configuration** NEW

55

Adds `sbx settings set kit.allowedSources` to configure the kit source allowlist, enabling Git kits to be pulled from sources beyond the default Docker Hub allowlist.

— Names the config command but no example or scope detail product docs

12 **Registry mirror redirection for sandbox images** NEW

53

Adds `platform.images.registryMirror` setting to redirect Docker Hub-resolving sandbox template and kit images to an organization's registry mirror.

— Config key named, no example or scope beyond one sentence v0.39.0

13 **Expanded sbx diagnose checks** IMPROVED

51

Extends `sbx diagnose` to report free disk space on the volume holding sandbox data and to detect broken, shadowed, or stale SSH client configuration.

— Names command and two new checks, no example output v0.39.0

14 **Declarative credential file structure for OAuth kits** NEW 48

Allows OAuth kits to declare their credential file using the declarative `credentialFile.structure` form, which is rendered to well-formed JSON.

— Field name given, no example of usage v0.39.0

15 **Remote control for Claude Code sandboxes** NEW 45

Enables Claude Code's `/remote-control` command inside a sandbox via the `claude.remoteControl` setting, toggled with `sbx settings set claude.remoteControl true`.

— Exact setting named but no worked example given v0.39.0

16 **Ubuntu 25.10 packages discontinued** BREAKING 30

Ubuntu 25.10 packages are no longer published because Ubuntu 25.10 reached end-of-life.

— States the removal with no migration guidance v0.39.0

└──▷ **BREAKING ON UPGRADE**

! After upgrading Docker Sandboxes, the NVIDIA kernel modules in the bundle must be rebuilt to match the new guest kernel — download the new release's `nerdbox-nvidia-modules-x86_64.zip` and re-run `prepare-nvidia-bundle.sh`.

! `sbx run --model` no longer accepts the `ollama/` prefix to select the Ollama backend; use the new `--provider ollama` flag instead.

! The `SANDBOX_VM_ID` environment variable inside sandboxes is deprecated in favor of `SANDBOX_NAME` and `SANDBOX_ID`.

! Ubuntu 25.10 packages are no longer published (Ubuntu 25.10 is end-of-life).

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SOAR

Tines

Feedback ⓘ 2 RELEASES · 2026-08-19 BLOG ↗

Tines is a no-code automation platform for security teams to build workflows that connect tools and respond to threats.

> Tines shipped a first-party MCP server that lets AI clients like Claude, Cursor and ChatGPT author Stories directly, alongside a new Workbench AI chat interface extended to Cases, dashboard templates with scheduled snapshot emails, built-in Apps functions for reading Cases/Records/Resources, SSO-gated access for Apps and Collections, customizable AI credit alerts, and new case list layouts.

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

01 **MCP server endpoint for AI story authoring** NEW 70

Tines Stories now ships a first-party MCP server endpoint built into every tenant, letting AI clients such as ChatGPT, Codex, Cursor, Claude Code and Claude Desktop create and modify stories through natural-language instructions against the same story-authoring capabilities as the storyboard UI. This built-in server is distinct from custom MCP servers built inside Tines Stories for running production integrations, and both can be used simultaneously.

— Names AI clients but gives no endpoint path or auth detail. product docs

02 **Workbench AI chat interface for Stories and Cases** NEW 70

Workbench is a new AI chat interface that lets practitioners ask questions and take action in real time using Tines Stories-powered tools, keeping data private within the tenant. Tool configurations, chats, and custom instructions are scoped per user, who can enable credentials or stories they have access to, and individual template actions can be turned on for use inside Workbench; a built-in formula documentation tool answers questions about Tines Stories formula usage and construction. Workbench has also been extended into Workbench for Cases, integrating Cases with AI.



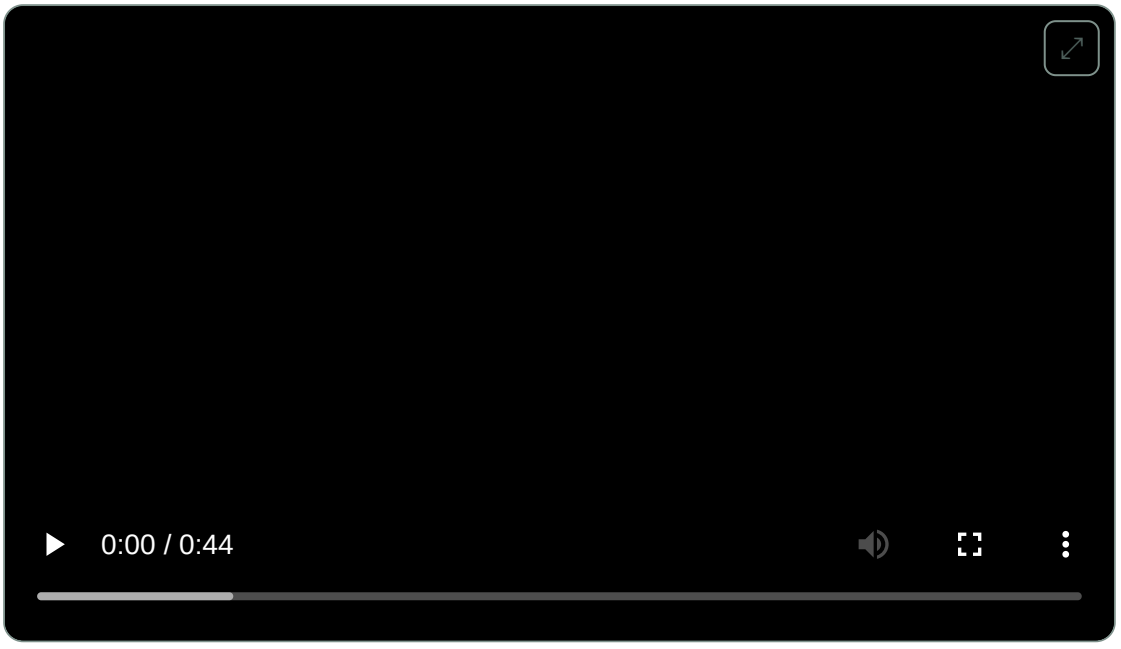
— Mechanism is described in detail but no config keys or API surfaces are named.

product docs

03 **Dashboard templates and scheduled snapshot emails** NEW

70

Tines Stories adds sample dashboard templates so practitioners can create a pre-built dashboard ready to rename or edit immediately, and adds dashboard Snapshots to email a copy of a dashboard on a configurable schedule, with control over recipients, frequency, and which charts to include — enabling recurring reports for stakeholders without Tines Stories access.



— Clear configurable options but no exact UI path or API named. product docs

04 **Built-in functions to read Cases, Records, Resources in Apps**

NEW

65

Tines Apps can now call built-in functions to read cases, records, and resources directly from your tenant, eliminating the need to create or call an endpoint — adding native support for Tines primitives inside Apps.

– Names the primitives but no function signatures or syntax given.

Support for Tines primitives in Apps (2…

THINNER COVERAGE BELOW

05 **SSO-based sharing and access for Apps and Collections** NEW

45

Tines 42.7.0 adds SSO-based sharing and access for Apps, allowing practitioners to gate automation apps behind their identity provider, and extends the same identity-provider-controlled access to Collections.

– No specific identity providers or config keys named. changelog-20260819-63b50e2e

06 **Customizable AI credit usage alerts** IMPROVED

35

Tines 42.7.0 enables customization of AI credit usage alerts so teams can tune notification thresholds for their consumption.

– No specific threshold values or configuration location named.

changelog-20260819-63b50e2e

07 **New case list layout options** IMPROVED

35

Tines 42.7.0 introduces new layout options for the case list view, giving analysts more flexibility in how they triage and review cases.

– Describes the change but not the specific layout options available.

changelog-20260819-63b50e2e

08 **More flexible events panel** IMPROVED

15

Tines Stories delivers a more flexible events panel.

– Only a one-line description with no mechanism or scope. product docs

WAS THIS USEFUL?  

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 3 RELEASES · SEEN 2026-08-19 NOTES ↗

> Iru expanded compliance coverage with a new Google Cloud Platform connector and 209 SSO application templates, alongside a small filtering improvement for unassigned actions.

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

01 **Google Cloud Platform compliance source** NEW 75

Adds Google Cloud Platform to the Iru Compliance connector catalog, enabling configuration and inventory evidence collection across Google Cloud IAM, Cloud Storage, Cloud KMS, and Cloud Logging. Connects to a Google Cloud project via a read-only service account and JSON key for scoped, least-privilege access.

– Names services and auth mechanism but no config steps

Google Cloud Platform is now available …

02 **209 SSO application templates** NEW 65

Adds 209 pre-configuration templates for SAML and OIDC single sign-on across commonly used web applications, including Slack, Zoom, and Google Workspace. Templates also cover SCIM user provisioning configuration, guiding identity administrators through the full setup process.

– Concrete count and protocols but no exact navigation

Single Sign On is a Few Clicks Away wit…

THINNER COVERAGE BELOW

03 **Unassigned filter for action owners** IMPROVED 50

Adds an Unassigned option to the Owner and Delegate filters on the Actions list in Compliance Automation, letting admins surface ownerless actions without manually scanning the table.

– Names exact UI filter but small scope Filter actions by unassigned owner or d…

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

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

> Trend Vision One extended Application Control policies to macOS agents, bringing feature parity with Windows and Linux including Lockdown mode support.

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

01 **Application Control support for macOS agents** NEW 62

Trend Vision One now supports Application Control for macOS agents in Endpoint Security Policies, enabling the same allow/block application rules previously available for Windows and Linux agents, including Lockdown mode. It is configurable under Endpoint Security → Endpoint Security Configuration → Endpoint Security Policies in the Vision One console.

— Names exact UI navigation path but no CLI/API or limits given.

TrendAI Vision One agents support Appli…

WAS THIS USEFUL? 👍 👎

Huntress

Feedback

ⓘ

1 RELEASE · 2026-08-19

CHANGELOG ↗

Huntress is a managed security platform that provides threat detection, incident response, and endpoint protection for small and mid-sized businesses.

> Huntress introduced a new Investigations dashboard giving SOC investigation transparency with timelines and exportable reports.

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

01 **Investigations dashboard with export** NEW 68

A new Investigations dashboard, accessed via Investigations in the top navigation, shows a chronological timeline of SOC signals reviewed, actions taken, remediations applied, and final outcomes for both closed-benign and reported investigations. Investigations can be exported as CSV or Excel for compliance reporting or post-incident records via Investigations > locate the investigation > Export.

Export a closed investigation for a compliance report or post-incident record.

> 📍 In the Huntress portal, go to Investigations (top navigation) › locate the investigation › click Export › choose CSV or Excel.

— Names UI path and export formats but no API or config. changelog-20260819-7ab40866

WAS THIS USEFUL? 👍 👎

Automox

Feedback

1 RELEASE · 2026-08-19

BLOG ↗

Automox is a cloud-based platform that automates patch management and security remediation across endpoints and servers.

> Automox added background access to Windows and macOS system utilities from Remote Tools, plus organization-level controls to enforce and lock session recording for audit compliance.

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

01 **System Tools access in Remote Tools** NEW 73

A new System Tools option under the Remote Tools button on Device Details opens Task Manager, Registry Editor, Device Manager, and Service Manager on managed Windows devices in a background session with no desktop takeover required. On macOS, System Tools surfaces Activity Monitor as the native equivalent of Task Manager for background process triage.

Restart a stalled Windows service on an end user's device without taking over their screen.

📍 In the Automox console, open Device Details for the target device › click Remote Tools › select System Tools — Service Manager opens in a new tab, running in the background.

— Names exact tools and UI path but only a UI flow, not an API/command.

Remote Tools: System Tools & Screen Rec…

02 **Screen recording controls for Remote Control sessions** NEW 60

A new Auto-record all sessions setting on the Remote Tools Settings tab automatically begins recording every Remote Control session at the organization level, and a new Lock Recording setting prevents technicians from stopping a recording mid-session to support audit and compliance requirements.

Ensure every Remote Control session is recorded and cannot be stopped by technicians, to satisfy audit requirements.

📍 In the Automox console, go to Remote Tools Settings tab › enable Auto-record all sessions › enable Lock Recording .

— Names both settings and their tab, but only UI steps given.

Remote Tools: System Tools & Screen Rec…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Managed Detection & Response (MDR)

Arctic Wolf

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

Arctic Wolf is a cloud-native security operations platform providing managed detection, response, and threat hunting services.

> Arctic Wolf introduced Data Explorer, giving customers direct query access to raw logs and analyzed platform data for security investigations.

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

01 **Data Explorer for raw and analyzed logs** NEW 31

Arctic Wolf Data Explorer lets customers query raw logs ingested by the Arctic Wolf platform to support incident investigations and answer security questions, alongside access to analyzed, processed platform data for deeper operational context.

— Describes purpose but no concrete query mechanism, UI path, or API named.

Arctic Wolf Data Explorer (2026-08-18)

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl

Feedback ⓘ 2 RELEASES · 2026-08-18 → 2026-08-19 BLOG ↗

Cribl is a data processing and routing platform that collects, transforms, and delivers security and observability data from multiple sources.

> Cribl shipped AWS Cloud Connections for automated multi-account log discovery, launched the Cribl Marketplace for in-product app discovery and installation, and added LlamaIndex support to LLM Observability.

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

01 **AWS Cloud Connections for automated log discovery** NEW 85

Cribl Stream can now automatically enumerate log destinations (S3, CloudWatch, Firehose, and Firehose destinations) across multiple AWS accounts and regions without manual hunting or custom scripts. It generates CloudFormation templates to provision a read-only AWS STS role for discovery via AWS APIs, and produces CloudFormation configuration files deployable as a stack to onboard discovered sources into Cribl Stream; results indicate whether a source is already onboarded, and the feature supports both existing Cribl organizations and brand-new AWS environments or accounts.

— Names mechanism, source types, and deployable CloudFormation artifact

Introducing AWS Cloud Connections: Auto…

02 **Cribl Marketplace app discovery hub** NEW 70

Cribl Marketplace is a new central in-product hub within Cribl.Cloud (for cloud and hybrid leaders) for browsing, discovering, and one-click installing prebuilt apps from Cribl, partners, and community developers, replacing manual GitHub URL imports and endpoint configuration. The catalog can be filtered by solution type, publisher, product, tags, support model, certification status, and recency, and per-app detail pages show overview, permissions, external API access, files, support information, and version history before install. Cribl plans a standardized publisher process, automated readiness checks, security scanning, review workflows, publisher verification, and a dedicated public Marketplace URL for non-customers.

— Details filters and app pages but roadmap items are unshipped

Introducing Cribl Marketplace: A centra…

THINNER COVERAGE BELOW

03 **LlamaIndex support in LLM Observability** NEW 41

Cribl LLM Observability now supports instrumenting LlamaIndex applications with OpenTelemetry to send telemetry to Cribl Stream and Cribl Search.

— Names the integration and data destinations but no setup steps `product docs`

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

1 RELEASE · 2026-08-19 BLOG ↗

Feedly Threat Intelligence aggregates security feeds and threat data to help organizations monitor emerging cybersecurity threats and vulnerabilities.

> Feedly Threat Intelligence launched a Suricata Rules AI Model that automatically harvests and exports detection rules from open-source reporting, while expanding its detection-rule extraction to cover six major rule formats.

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

01 Suricata Rules AI Model with expanded rule-format coverage

NEW 81

A new Suricata Rules AI Model can be paired with an AI Feed to continuously collect articles containing Suricata rules from 5,000 vetted CTI sources and repositories as new reporting is published, extract the rules, and make them exportable or pullable via API directly into detection pipelines, IDS, or SIEM tooling. Feedly Threat Intelligence now recognizes six detection rule formats across open-source reporting: Suricata, Snort, Sigma, YARA, Splunk, and KQL-based Hunting Queries for Microsoft Defender and Sentinel, spanning network traffic inspection, file/malware matching, and log-based detection.

Continuously surface new Suricata rules relevant to a specific threat actor or malware family as open-source reporting is published.

📍 In Feedly Threat Intelligence, go to AI Feeds › Create AI Feed, add the Suricata Rules AI Model, then pair it with additional AI Models representing your intelligence requirements (e.g. a specific threat actor or malware family) to filter incoming articles to only those containing matching Suricata rules.

— Names AI Model, source count, formats, API export; UI path given but no exact command

Continuously collect Suricata rules mat…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

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

> Tenzir Test 2.0.0 improves integration test diagnostics by capturing node stdout/stderr into log files, reporting node crash details directly in the harness, and making diagnostic output deterministic in test baselines (a breaking change for existing baselines).

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

01 **Node output capture in test fixtures** NEW 90

Test runs now capture `tenzir-node` output from startup through teardown into `node-stdout.log` and `node-stderr.log`, exposed to tests via the `TENZIR_NODE_STDOUT_LOG` and `TENZIR_NODE_STDERR_LOG` environment variables. A new `--keep` flag retains these per-test log files in the scratch directory after a test run instead of deleting them.

Inspect node stderr output during a test run to diagnose startup or crash issues.

```
$ tail "$TENZIR_NODE_STDERR_LOG"
```

— Names log files, env vars, flag, and includes a runnable command product docs

THINNER COVERAGE BELOW

02 **Deterministic diagnostic ordering in test baselines** BREAKING 55

TQL test baselines now record diagnostics before normal pipeline output, making baseline ordering deterministic. Existing baselines that contain both diagnostics and pipeline output channels must be updated to reflect the new ordering.

— Explains before/after and gives a migration note

Tenzir Test v2.0.0: Deterministic test …

03 **Node crash diagnostics in test harness** IMPROVED 45

When the node exits while its fixture is active, the harness now reports the exit code and a tail of stderr instead of leaving dependent tests to fail with opaque connection errors.

— Describes behavior change but no config or command surface product docs

└──▷ **BREAKING ON UPGRADE**

! TQL test baselines now record diagnostics before normal pipeline output; existing baselines that contain both diagnostics and pipeline output must be updated.

! Test baselines that contain both diagnostics and pipeline output channels must be updated, as diagnostics are now recorded before pipeline output.

WAS THIS USEFUL? 👍 👎

◆ **SIEM**

Sumo Logic

Feedback ⓘ **1 RELEASE · 2026-08-19** **BLOG ↗**

Sumo Logic is a cloud-based monitoring and analytics platform for logs, metrics, and security data.

> Sumo Logic's Scheduled Searches now support multiple independent notification channels firing from a single alert trigger.

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

01 **Multi-notification support for Scheduled Searches** **IMPROVED**

55

Scheduled Searches now support up to 10 notifications per search, each configured independently, all sharing the same trigger condition and firing from the same search execution.

— Explains mechanism and limit but no UI/API path given

Multi-Notification Support for Schedule…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder

Feedback ⓘ 1 RELEASE · 2026-08-19 CHANGELOG ↗

Intruder is a cloud-based vulnerability scanner that automatically discovers and tests systems for security weaknesses.

> Intruder shipped a UI overhaul in this window, adding a dark theme and redesigning navigation for faster, single-click access with pinnable favourites.

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

01 **Redesigned single-click navigation with Favourites** IMPROVED

55

Navigation was redesigned to give single-click access to every page, with pages grouped into expandable theme sections and icon/pill indicators flagging new data or activity. A pinnable Favourites shortcut bar was also added to the sidebar for quick access to most-visited pages.

– Describes UI mechanism and location but no exact paths changelog-20260819-85d92648

02 **Dark theme** NEW

18

Intruder added a dark theme option, per the release summary.

– Only mentioned in summary, no detail on how to enable changelog-20260819-85d92648

WAS THIS USEFUL? 👍 👎

runZero

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

runZero is an asset discovery and inventory platform that identifies and maps all connected devices across network infrastructure.

> runZero shipped granular self-hosted licensing kill switches via environment variables and added OCSF 1.9.0 export for asset inventory data.

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

01 Per-entitlement kill switches for self-hosted instances NEW

90

Self-hosted runZero instances can force-disable any license entitlement instance-wide via `RUNZERO_DISABLE_<ENTITLEMENT>` environment variables, overriding the license for every account and organization. Supported variables include `RUNZERO_DISABLE_AUTOSCAN`, `RUNZERO_DISABLE_AI`, `RUNZERO_DISABLE_PASSIVE_MONITORING`, `RUNZERO_DISABLE_API_ORG`, `RUNZERO_DISABLE_API_ACCOUNT`, `RUNZERO_DISABLE_DATA_RETENTION`, `RUNZERO_DISABLE_OFFLINE`, `RUNZERO_DISABLE_METAMIN`, `RUNZERO_DISABLE_FINDINGS`, `RUNZERO_DISABLE_ISSUES`, `RUNZERO_DISABLE_SUPPRESSION`, `RUNZERO_DISABLE_RISK_EVENTS`, `RUNZERO_DISABLE_RULES_BASIC`, `RUNZERO_DISABLE_REPORTS_BASIC`, `RUNZERO_DISABLE_REPORTS_ADVANCED`, `RUNZERO_DISABLE_CERTIFICATE_INVENTORY`, and integration-specific variables such as `RUNZERO_DISABLE_INT_CROWDSTRIKE`, `RUNZERO_DISABLE_INT_SENTINELONE`, `RUNZERO_DISABLE_INT_TENABLE`, `RUNZERO_DISABLE_INT_SPLUNK`, `RUNZERO_DISABLE_INT_SERVICENOW`, `RUNZERO_DISABLE_INT_QUALYS`, `RUNZERO_DISABLE_INT_INSIGHTVM`, `RUNZERO_DISABLE_INT_CLOUD`, `RUNZERO_DISABLE_INT_LDAP`, `RUNZERO_DISABLE_INT_AZURE_AD`, `RUNZERO_DISABLE_INT_INTUNE`, `RUNZERO_DISABLE_INT_WIZ`, `RUNZERO_DISABLE_INT_JIRA`, `RUNZERO_DISABLE_INT_PAN`, and more.

Disable autonomous (autoscan) discovery and AI features across an entire self-hosted runZero instance before a maintenance window.

```
$ RUNZERO_DISABLE_AUTOSCAN=true
RUNZERO_DISABLE_AI=true
```

Lock down all third-party integration imports (e.g. CrowdStrike, SentinelOne, Tenable) on a restricted self-hosted deployment.

```
$ RUNZERO_DISABLE_INT_CROWDSTRIKE=true
RUNZERO_DISABLE_INT_SENTINELONE=true
RUNZERO_DISABLE_INT_TENABLE=true
RUNZERO_DISABLE_INT_QUALYS=true
RUNZERO_DISABLE_INT_INSIGHTVM=true
```

— Names every env var and shows exact usage, missing only rollout guidance. product docs

02 **OCSF 1.9.0 export for asset inventory** NEW 74

Adds OCSF 1.9.0 export for the asset inventory, producing `assets.ocsf.jsonl` and `assets.ocsf.json` files. Assets export as Device Inventory Info events and their observed services as Live Evidence Info events; fields without an OCSF equivalent are preserved under the event's `unmapped` object so no asset data is silently dropped. The export is accessible via the inventory export menu and the Export API, supporting both `.jsonl` and `.json` formats.

— Names formats, schema events and API, but no endpoint or CLI syntax given.

runZero 5.1.260818.1 (2026-08-18)

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Qualys Patch Management

Feedback ⓘ 1 RELEASE · 2026-08-18

BLOG ↗

Qualys Patch Management is a cloud-based tool that identifies, prioritizes, and automates the deployment of security patches across IT infrastructure.

> Qualys Patch Management introduces AI-powered scoring to predict patch deployment success before rollout, and improves remediation state reporting for KEV compliance tracking.

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

01 **AI-Powered Patch Reliability Scoring** NEW 60

A new Qualys AI-Powered Patch Reliability Scoring feature predicts patch deployment success before rollout by combining LLM-analyzed global public sentiment (technical discussions, release feedback) with Qualys telemetry on patch rollback rates and vulnerability reopen rates into a single reliability score.

.
.
.

— Explains mechanism and inputs but no UI path or config surface 2026-08-18-7d12e417

THINNER COVERAGE BELOW

02 **TruRisk Eliminate true remediation state reporting** IMPROVED 45

TruRisk Eliminate now reports true remediation state — including reboot completion — so teams can verify KEV deadline compliance after deploying cumulative updates.

— States what changed but no concrete steps or config 2026-08-18-7d12e417

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

authentik

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

The authentication glue you need.

> authentik's headline addition this window is AKQL, a structured query language for searching events, users and groups across the Admin UI and REST API, alongside a wave of Enterprise features (Workspace/Entra ID sync, RADIUS EAP-TLS, Account Lockdown, Object Lifecycle Management) and infrastructure work spanning bootstrapped admin accounts, PostgreSQL read-replica/direct-connection routing, a one-click AWS CloudFormation deployment, a theme-aware centralized file manager, and proxy-provider host/backend overrides.

└──▷ WHAT SHIPPED · 26 FEATURES most completely described first ? what's the number?

01 AKQL structured query language for search and policies NEW

100

AKQL (authentik Query Language), built on DjangoQL, adds structured search for events, users, and groups in the Admin interface, e.g. `action = "login"` or `context.geo.country = "Germany"`, with dot-notation into nested fields like `context.geo.country`, `context.http_request.path`, and `attributes.department`. It is exposed via the `search` query parameter on REST API list endpoints, e.g. `GET /api/v3/core/users/?search=username%20%3D%20%22bob%22`, and via a new `Query` field on Event Matcher policies, combinable with existing fields like `Action` or `App`. Supported operators are `=`, `!=`, `>`, `>=`, `<`, `<=`, `~`, `!~`, `startswith`, `not startswith`, `endswith`, `not endswith`, in (...), and not in (...), over typed strings, integers, floats, booleans, null, and lists. Searchable fields include Events' `action`, `event_uuid`, `app`, `client_ip`, `user.username`, `user.email`, `brand.name`, `context.*`, `created`; Users' `username`, `name`, `email`, `path`, `is_active`, `type`, `attributes.*`; and Groups' `name`, `is_superuser`, `attributes.*`. Invalid input falls back to plain case-insensitive substring search, with autocomplete via `↓` / `↑`, `Tab` / `Enter`, and `ESC`.

Query the REST API for users with a specific username using AKQL via the `search` parameter.

```
$ curl -s 'https://authentik.example.com/api/v3/core/users/?search=username%20%3D%20%22bob%22' -H 'Authorization: Bearer <token>'
```

Restrict an Event Matcher policy to flag failed logins originating outside the United States using the `Query` field.

> 📍 In the Admin interface, go to Customization › Policies, create or edit an Event Matcher Policy, and enter in the Query field: action = `login_failed` and `context.geo.country` != `United States`

— Names the endpoint, field, operators and searchable fields exhaustively. `product docs`

---

02 **PostgreSQL read replicas, direct routing and TLS config**

`IMPROVED` 100

Adds indexed read replica config via `AUTHENTIK_POSTGRESQL__READ_REPLICAS__<n>__HOST/NAME/USER/PORT/PASSWORD/SSLMODE/SSLROOTCERT/SSLCERT/SSLKEY/CONN_MAX_AGE/CONN_HEALTH_CHECKS/DISABLE_SERVER_SIDE_CURSORS` with automatic read/write routing to the primary, and a direct/session-mode bypass via `AUTHENTIK_POSTGRESQL__DIRECT__HOST`, `AUTHENTIK_POSTGRESQL__DIRECT__PORT`, `AUTHENTIK_POSTGRESQL__DIRECT__NAME`, `AUTHENTIK_POSTGRESQL__DIRECT__USER`, `AUTHENTIK_POSTGRESQL__DIRECT__PASSWORD`, `AUTHENTIK_POSTGRESQL__DIRECT__SSLMODE`, `AUTHENTIK_POSTGRESQL__DIRECT__SSLROOTCERT`, `AUTHENTIK_POSTGRESQL__DIRECT__SSLCERT`, and `AUTHENTIK_POSTGRESQL__DIRECT__SSLKEY` to keep LISTEN/NOTIFY and advisory-lock traffic (Channels, dramatiq, migrations) off a transaction-pooling pooler. Adds primary-connection TLS via `AUTHENTIK_POSTGRESQL__SSLMODE` (default `verify-ca`), `AUTHENTIK_POSTGRESQL__SSLROOTCERT`, `AUTHENTIK_POSTGRESQL__SSLCERT`, `AUTHENTIK_POSTGRESQL__SSLKEY`; `AUTHENTIK_POSTGRESQL__CONN_HEALTH_CHECKS` (default `false`); `AUTHENTIK_POSTGRESQL__DISABLE_SERVER_SIDE_CURSORS` (default `false`); `AUTHENTIK_POSTGRESQL__CONN_MAX_AGE` (`0` default closes per-request, positive integer sets lifetime, `None` unlimited); and `AUTHENTIK_POSTGRESQL__DEFAULT_SCHEMA` (default `public`, must be set pre-first-startup). Also adds `env://<name>?<default>` / `file://<name>?<default>` URI loading for config values, the `ak dump_config` command, and confirms `AUTHENTIK_POSTGRESQL__HOST/PORT/USER/PASSWORD` hot-reload without a restart.

Route LISTEN/NOTIFY and advisory-lock traffic away from a PgBouncer transaction-pool pooler to avoid silent breakage of Channels and dramatiq workers.

```
$ AUTHENTIK_POSTGRESQL__HOST=pgbouncer
AUTHENTIK_POSTGRESQL__PORT=5432
AUTHENTIK_POSTGRESQL__DISABLE_SERVER_SIDE_CURSORS=true
AUTHENTIK_POSTGRESQL__DIRECT__HOST=postgres-primary
AUTHENTIK_POSTGRESQL__DIRECT__PORT=5432
```

Add a PostgreSQL read replica with mutual TLS to distribute query load away from the primary.

```
$
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__HOST=replica.db.internal
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__PORT=5432
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__NAME=authentik
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__USER=authentik
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__PASSWORD=file:///run/secr
ets/db_password
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__SSLMODE=verify-full
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__SSLROOTCERT=/certs/ca.pem
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__SSLCERT=/certs/client.pem
AUTHENTIK_POSTGRESQL__READ_REPLICAS__0__SSLKEY=/certs/client.key
```

Verify the full resolved authentik configuration after applying environment variable changes in Docker Compose.

```
$ docker compose run --rm worker ak dump_config
```

— Enumerates every env var, default and the verification command. product docs

03 **Bootstrap akadmin account and hashed-password APIs** NEW

95

Adds `AUTHENTIK_BOOTSTRAP_PASSWORD_HASH` (pre-hashed, generated via `docker compose run --rm server hash_password`), `AUTHENTIK_BOOTSTRAP_PASSWORD` (plaintext; combining both with the hash var errors), `AUTHENTIK_BOOTSTRAP_TOKEN`, and `AUTHENTIK_BOOTSTRAP_EMAIL` environment variables to configure the `akadmin` user at first startup, bypassing the out-of-box experience. The same values are settable in Helm via `bootstrap_password_hash`, `bootstrap_token`, and `bootstrap_email` under `authentik:` or a Kubernetes secret via `global.envFrom.secretRef`. Post-install, a new `POST /api/v3/core/users/<id>/set_password_hash/` endpoint (requires `authentik_core.reset_user_password` permission, hash in the `password` field)

sets hashed passwords on existing users or service accounts, and blueprints support the same behavior via a `password_hash` user attribute.

Generate a hashed password non-interactively before first deployment to populate `AUTHENTIK_BOOTSTRAP_PASSWORD_HASH` without exposing the password in shell history.

```
$ echo "$PASSWORD" | docker compose run --rm server hash_password
```

Bootstrap the akadmin account automatically in Docker Compose using a pre-hashed password stored safely in a `.env` file.

```
$
AUTHENTIK_BOOTSTRAP_PASSWORD_HASH='pbkdf2_sha256$1000000$xKKFuYtJ
EE27km09BD49x2$4+Z6j3utmouPF5mik0Z21L2P0og5IlmMmIJ46Tj3zCM='
AUTHENTIK_BOOTSTRAP_TOKEN=your-token-here
AUTHENTIK_BOOTSTRAP_EMAIL=admin@example.com
```

Set a hashed password on an existing user post-install via the API, scoped to a service account or regular user.

```
$ curl -X POST
https://authentik.example.com/api/v3/core/users/1/set_password_ha
sh/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"password":
"pbkdf2_sha256$1000000$xKKFuYtJEE27km09BD49x2$4+Z6j3utmouPF5mik0Z
21L2P0og5IlmMmIJ46Tj3zCM="}'
```

— Full env var, endpoint, permission and Helm key names with commands. product docs

---

04 Centralized file management and theme-aware asset picker NEW

90

A new `Customization > Files` section in the Admin interface provides centralized upload, delete, and search for branding and application-icon images, backed by the default `/data` directory or by S3 storage. File picker fields (Brand Logo, Brand Favicon, Brand Default flow background, Flow Background, Application Icon, Source Icon) accept uploaded files, built-in static assets, external URLs, Font Awesome icons via `fa://fa-circle-user -style` syntax (icon fields only), and theme-aware paths using a `%(theme)s` placeholder resolving to `light` or `dark`, including for S3 presigned URLs. Uploads are constrained to relative paths (no leading slash, `..`, `//`, or leading `.`), max 25 MB per file, max 1024-character full path, max 255 characters per segment, and

characters limited to letters, numbers, `.`, `_`, `-`, `/`, and `%(theme)s`. The Files API returns `themed_urls` maps, and Brand/entity APIs expose `branding_logo_themed_urls`, `branding_favicon_themed_urls`, `background_themed_urls`, `meta_icon_themed_urls`, and `icon_themed_urls` fields.

Upload a company logo and reference it as a theme-aware brand logo that serves different assets for light and dark mode.

> 📍 In the Admin interface, go to Customization › Files, upload your light and dark logo variants, then in your brand settings set the logo field to a `%(theme)s`-based path (e.g. `/media/branding/%(theme)s/logo.png`) so authentik resolves the correct variant per theme.

— Names UI path, placeholders, limits and every themed API field. product docs

---

05 **Global and per-stage SMTP email configuration** NEW

90

Adds global SMTP settings via `AUTHENTIK_EMAIL__HOST`, `AUTHENTIK_EMAIL__PORT`, `AUTHENTIK_EMAIL__USERNAME`, `AUTHENTIK_EMAIL__PASSWORD`, `AUTHENTIK_EMAIL__USE_TLS`, `AUTHENTIK_EMAIL__USE_SSL`, `AUTHENTIK_EMAIL__TIMEOUT`, and `AUTHENTIK_EMAIL__FROM` in `.env` (Docker) or the equivalent `email.host`, `email.port`, `email.username`, `email.password`, `email.use_tls`, `email.use_ssl`, `email.timeout`, `email.from` keys under `authentik:` in `values.yaml` (Helm). Adds the `ak test_email <to_address>` CLI command (with `-S <stage_name>` to target a specific Email stage), the `SSL_CERT_FILE` env var for a custom CA bundle, and support for per-Email-stage SMTP overrides.

Verify that global SMTP email delivery is working from the authentik worker container in a Docker Compose deployment.

```
$ docker compose exec worker ak test_email [email protected]
```

Test a specific Email stage's SMTP settings (rather than the global config) to confirm stage-level overrides are routing correctly.

```
$ docker compose exec worker ak test_email [email protected] -S
my-email-stage
```

— Names every env var, CLI flag and Helm key with runnable commands. product docs

---

06 **Proxy outpost header injection and sign-out endpoint** IMPROVED

80

Documents the `additionalHeaders` attribute on groups or users for injecting static upstream headers (e.g. `X-test-header: test-value`), the full set of headers the proxy outpost sets — `X-authentik-username`, `X-authentik-groups`, `X-authentik-entitlements`, `X-authentik-email`, `X-authentik-name`, `X-authentik-uid`, `X-authentik-meta-outpost`, `X-authentik-meta-provider`, `X-authentik-meta-app`, `X-authentik-meta-version`, and `X-Forwarded-Host` (proxy mode only) — and the `/outpost.goauthentik.io/sign_out` logout endpoint, available on the protected-application host in proxy/single-app mode or the auth URL in domain-level mode, which invalidates the session scoped to that outpost.

— Names every header, the attribute and the endpoint path. product docs

---

07 **Dynamic host header and backend override for proxy provider**

NEW

75

The `ak_proxy` scope-mapping expression gains a `host_header` key (authentik 2025.6.1+) to dynamically override the host header sent upstream, and a `backend_override` key to dynamically select the upstream backend per user in proxy mode.

Dynamically route each user to a per-user upstream backend in proxy mode — useful when different users need isolated backend instances.

python

```
return {
    "ak_proxy": {
        "backend_override":
f"http://foo.bar.baz/{request.user.username}"
    }
}
```

Override the host header sent to the upstream to match the provider's configured Internal host — prevents host-mismatch rejections on strict internal services (authentik 2025.6.1+).

python

```
from urllib.parse import urlparse
parsed_url = urlparse(provider.proxyprovider.internal_host)
return {
    "ak_proxy": {
        "host_header": parsed_url.netloc
    }
}
```

— Names both scope-mapping keys with runnable Python expressions. product docs

---

08 **One-click AWS deployment via CloudFormation** NEW 75

A new CloudFormation template deploys a production-ready authentik stack: VPC, RDS PostgreSQL Multi-AZ cluster, ECS cluster running server and worker tasks, an ALB with an ACM certificate, an EFS filesystem, and AWS SSM-stored credentials. The stack takes a Certificate ARN input from AWS Certificate Manager and outputs the ALB endpoint; initial setup is reached at `http://<domain_you_configured>` and prompts for the `akadmin` password.

— Names every provisioned resource and the setup URL, no CLI given. product docs

09 **Unauthenticated Paths and URLs on proxy provider** NEW 70

New Unauthenticated Paths and Unauthenticated URLs fields on the proxy provider allow unauthenticated access to specific paths or URLs, defined per-line as Rust `regex` crate expressions; in proxy and single-application forward-auth modes expressions match the request path, while in domain-level forward-auth mode they match the full URL including scheme and host. Domain-level forward auth cannot enforce per-application authorization rules, so single-application mode is required when each application needs its own policies or bindings.

— Names the fields and regex engine but gives no example expression. product docs

THINNER COVERAGE BELOW

10 **RADIUS provider EAP-TLS and message-authenticator handling** IMPROVED 50

The RADIUS provider adds EAP-TLS support, authenticating network clients using client certificates. It also now allows an empty message authenticator, broadening compatibility with RADIUS clients that omit this field.

— Names the protocol changes but no config flag or field. version/2026.8.0

11 **Client certificate authentication for device login** NEW 50

Users can enroll or authenticate via client certificates issued to devices, smart cards, PIV cards, or hardware tokens.

— Names supported credential types but no setup steps. product docs

12 **Account Lockdown for compromised accounts** NEW 50

Account Lockdown immediately secures a compromised account by disabling it, revoking its tokens, ending its sessions, and recording the action in the audit log.

— Explains the mechanism but no UI path or API given. product docs

13 **Object Lifecycle Management for periodic reviews** NEW

50

A preview Object Lifecycle Management feature schedules periodic reviews of applications, groups, and roles, assigns reviewers, and tracks overdue reviews.

— Explains scheduling/reviewer mechanism but preview with no config. product docs

14 **Device trust connectors for Fleet and Chrome** NEW

50

A new Fleet connector uses device information from Fleet, and a new Google Chrome connector uses Chrome Enterprise Device Trust signals, in device-aware access decisions.

— Names both connectors but no config steps. product docs

15 **Shared Signals Framework and WS-Federation providers** NEW

45

A new Shared Signals Framework provider sends security events to subscribed applications, including Apple Business Manager. A new WS-Federation provider connects WS-Federation applications to authentik for single sign-on.

— Names both protocols and one consumer app, no config detail. product docs

16 **OAuth authentication for SCIM provisioning** IMPROVED

45

SCIM provisioning now authenticates using short-lived OAuth tokens instead of static tokens.

— States the before/after but no config or endpoint named. product docs

17 **External OAuth and SAML sources embedded in flows** NEW

45

External OAuth and SAML sources can now be embedded directly in a flow, enabling dynamic redirection to an external identity provider such as Entra ID or Google Workspace before the flow continues.

— Describes the mechanism but no flow-stage name or config field. product docs

18 **Session and OAuth2 logout cleanup on deactivation** IMPROVED

45

Deactivating a user now sends back-channel logout requests to OAuth2 providers to terminate downstream sessions, and immediately deletes all of that user's authentik sessions.

— Describes the mechanism but no config toggle or endpoint. version/2026.8.0

19 **Deprecated PostgreSQL connection-options variables** DEPRECATED

45

`AUTHENTIK_POSTGRESQL__CONN_OPTIONS` and `AUTHENTIK_POSTGRESQL__REPLICA_CONN_OPTIONS` are deprecated as of version 2026.5 and will be removed in a future version.

— Names the deprecated vars but gives no migration path. product docs

20 **Identity provisioning sync to Google Workspace and Entra ID** NEW 40

New Enterprise integrations synchronize users and groups from authentik outward to Google Workspace and to Microsoft Entra ID.

— Names the two sync targets but no config surface or mechanism. product docs

21 **Event maps, charts and CSV exports** NEW 40

New event maps and charts visualize recent events by location, type, and volume, and new CSV data exports allow export of user and event data for analysis, reporting, or backup.

— Names two reporting additions with no export path or endpoint. product docs

22 **Multiple simultaneous user account sessions** NEW 40

Multiple user accounts can now remain signed in within the same browser simultaneously, with switching available from the User interface header.

— Names the UI location but no config or session limit. product docs

23 **SAML audience override and source stage failure action** NEW 40

SAML sources gain an audience override field to specify a custom SAML audience value, and Enterprise source stages gain a configurable failure action controlling what happens when a source stage fails.

— Names both new fields with no exact key or UI path. version/2026.8.0

24 **Enhanced audit logging with before/after values** IMPROVED 35

Audit logging now records detailed object changes, including before-and-after values, for compliance review.

— States the before/after mechanism but no surface to view it. product docs

25 **Local device login for Windows and Linux** NEW 30

A preview Local device login feature enables users to sign in to Windows and Linux devices using their authentik credentials.

— One-line preview description, no setup detail. product docs

26 **Query string preserved in proxy outpost redirects** IMPROVED

30

Post-authentication redirects for the proxy outpost now include the original query string, preserving URL parameters through the login flow.

— One-line behavior fix with no config or example. version/2026.8.0

└──▷ **BREAKING ON UPGRADE**

! `AUTHENTIK_POSTGRESQL__CONN_OPTIONS` is deprecated as of version 2026.5 and will be removed in a future version.

! `AUTHENTIK_POSTGRESQL__REPLICA_CONN_OPTIONS` is deprecated as of version 2026.5 and will be removed in a future version.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

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

> Clerk added biometric sign-in support for returning users across its mobile SDKs.

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

01 **Biometric sign-in for mobile SDKs** NEW

45

Clerk enables biometric sign-in via a device-trusted prompt for returning users on Expo, iOS, and Android mobile SDKs.

— Names three platforms but no API, flag, or config detail.

Biometric sign-in for Expo, iOS, and An…

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

SailPoint is an identity governance and access management platform that automates user access provisioning, certification, and compliance across enterprise systems.

> SailPoint's Separation of Duty policies can now use Roles and Access Profiles as criteria, not just entitlements.

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

01 **Roles and Access Profiles as SoD criteria** NEW 62

SoD policy List A and List B criteria now support SailPoint Roles and Access Profiles in addition to entitlements, configurable via Admin → Policies → Separation of Duty, including mixed criteria across lists (e.g. entitlements in List A and a Role in List B). Existing conflict detection, access-request approval visibility, and violation management workflows continue to work for policies using Role or Access Profile criteria.

⋮

— Names exact UI path and behaviour but no API/config key.

New Capability: Roles Criteria in SoD P…

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback

1 RELEASE · SEEN 2026-08-19

NOTES ↗

Nudge Security monitors SaaS application usage and security risks across organizations to reduce cloud infrastructure threats.

> Nudge Security now detects AI agents your team builds on Datadog and HubSpot, adding them to its AI agent inventory alongside existing risk insights.

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

01 **AI agent discovery for Datadog and HubSpot** NEW 80

Nudge Security's browser extension now discovers AI agents built on Datadog and HubSpot, surfacing each one in the AI agents inventory under Identities with its creator, components, and risk insights. Discovery requires enabling browser-based agent discovery in the browser extension settings, and access is currently offered as part of an AI agent discovery research preview.

Enable browser-based AI agent discovery so Datadog- and HubSpot-built agents automatically appear in your inventory.

> 📍 In the Nudge Security browser extension, go to Settings and turn on browser-based agent discovery. Then navigate to Identities › AI Agents tab to view discovered agents and, if needed, request access to the research preview.

— Clear enable steps and inventory location, but only two integrations named

Discover AI agents built on Datadog and…

WAS THIS USEFUL? 👍 👎

1Password

Feedback

ⓘ

2 RELEASES · 2026-08-18 → 2026-08-19

CHANGELOG ↗

1Password is a password manager that securely stores, generates, and autofills login credentials and sensitive information across devices.

> 1Password's SCIM Bridge 2.9.16 hardens provisioning infrastructure with per-IP rate limiting and post-quantum TLS support.

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

01 **Rate limiting and post-quantum TLS in SCIM Bridge** IMPROVED

43

SCIM Bridge 2.9.16 adds per-IP rate limiting to all SCIM Bridge routes, protecting provisioning endpoints from abuse, and adds TLS support for post-quantum key exchange across SCIM Bridge connections.

— Names the mechanisms but no config keys or flags given.

1Password SCIM Bridge 2.9.16 (2026-08-1… changelog-20260819-0cb53b11

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ 1 RELEASE · 2026-08-19 BLOG ↗

Ping Identity is an enterprise identity and access management platform providing authentication, authorization, and single sign-on solutions.

> Ping Identity's PingOne added bidirectional Microsoft Entra ID provisioning along with Entra-based group attribute mapping and credential management from user profiles, while the PingOne Advanced Identity Cloud rapid channel shipped ESV support, navigation reorganization and a batch of usability and accessibility fixes.

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

01 **External Group Names mapping from Entra ID** NEW 70

Supports mapping attributes from Microsoft Entra ID to the External Group Names attribute in PingOne, enabling group-based application access control using Microsoft group membership.

— Names exact attribute and its access-control use case. product docs

02 **Reorganized navigation for identity and auth settings**

IMPROVED 60

Moves the Identity Verification page under Integrations, Android Key Attestation under Authentication Methods, and Security Questions under Authentication.

— Names exact pages and new menu locations for navigation. 18 Aug 2026 (2026-08-19)

THINNER COVERAGE BELOW

03 **Microsoft Entra ID provisioning connection** NEW 55

Adds a Microsoft Entra ID provisioning connection, enabling bidirectional provisioning of users and groups between Microsoft Entra ID and PingOne.

— Names the integration and sync direction but no setup steps. product docs

04 **ESV support in connector server config and script editor** NEW

55

Enables use of an ESV for the connector server name in the New Connector Server modal, and surfaces errors in the script editor when scripts reference environment ESVs incorrectly.

— Names both surfaces but not the ESV mechanism itself. 18 Aug 2026 (2026-08-19)

05 **Credential reset and delete from user profile** NEW

40

Adds reset and delete capabilities for credential records directly from a user's profile in PingOne.

— Clear UI location but no further mechanism described. product docs

06 **Continuous identity creation in creation modal** IMPROVED 40

Allows creating another identity immediately after creating one, without closing the modal.

— Simple workflow tweak, clearly described but minor. 18 Aug 2026 (2026-08-19)

07 **Accessibility and usability improvements** IMPROVED 40

Supports closing modals with the Escape key, improves keyboard navigation and accessibility information across interactive controls, and ensures forms clearly identify fields containing errors after submission.

— Groups several small UX fixes without deep mechanism detail. 18 Aug 2026 (2026-08-19)

08 **Editable start time for scheduled jobs** IMPROVED 25

Enables editing a future-scheduled job to start sooner.

— Very thin description with no scope or interface details. 18 Aug 2026 (2026-08-19)

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ 1 RELEASE · 2026-08-18

NOTES ↗

DigiCert CertCentral is a platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizations.

> CertCentral adds a reusable Persistent DNS TXT record method for bulk domain revalidation, extends X9 PKI for TLS certificates to 2- and 3-year terms, and allows in-place ICA updates on existing ACME Directory URLs.

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

01 In-place ICA update for ACME Directory URLs IMPROVED

85

Existing ACME Directory URLs can now have their issuing intermediate CA (ICA) certificate updated directly via Automation > ACME Directory URLs > Actions > Update ACME Directory URL, selecting a new 'Intermediate chains [Intermediate CA] > [Root CA]' option, replacing the prior requirement to create a brand-new ACME Directory URL to change the certificate chain.

Switch an existing ACME Directory URL to a different intermediate CA chain without creating a new URL — useful when rotating ICA certificates across automated issuance pipelines.

📍 In CertCentral, go to Automation › ACME Directory URLs. In the Actions menu for the target URL, select Update ACME Directory URL. Under ACME Directory URL Restrictions, use the 'Intermediate chains [Intermediate CA] › [Root CA]' menu to choose the desired chain, then select Update ACME Directory URL to save.

— Exact navigation path and menu names make this directly runnable August 18, 2026

02 Persistent DNS TXT record bulk revalidation NEW

80

Under Certificates > Domains, a new Revalidation by Persistent DNS TXT record option lets a one-time DNS TXT record containing a DigiCert-provided persistent URI be reused for every subsequent domain revalidation cycle, removing the need to create a new record or URI value each time.

Revalidate a batch of domains using the Persistent DNS TXT record method so the same DNS record can be reused on every future revalidation cycle.

📍 In CertCentral, go to Certificates › Domains. Select the domains to revalidate, open the Submit domains for revalidation menu, and choose Revalidation by

Persistent DNS TXT record . Follow the on-screen instructions to place the DigiCert-provided persistent URI in your DNS TXT record once, then submit.

— Names exact menu path and mechanism, no numeric limits given August 18, 2026

03 **2- and 3-year X9 PKI for TLS validity** IMPROVED 75

X9 PKI for TLS certificate validity options are extended from a maximum of 397 days to selectable 1-, 2-, or 3-year terms, with custom validity capped at 1,095 days, chosen in the Validity period section of the Request X9 PKI for TLS Certificate page.

— Concrete before/after limits and exact page name, but no command August 18, 2026

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ 2 RELEASES · 2026-08-18 → 2026-08-19

CHANGELOG ↗

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale added pagination to the list tailnets API, changing the default response to 100 results and requiring callers with larger organizations to page through results.

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

01 **Pagination for list tailnets API endpoint** BREAKING

90

The list tailnets API endpoint now supports `limit` and `cursor` query parameters, returning up to 100 results per page by default along with a `totalCount` field. The response includes a `cursor` value to pass as the `cursor` query parameter on subsequent requests; iteration ends when the returned `cursor` is empty. This is a breaking change: organizations with more than 100 tailnets previously received all results in a single response but must now paginate using `cursor` to retrieve the full set.

Page through all tailnets in a large organization by following the returned `cursor` until it is empty.

```
$ curl -H 'Authorization: Bearer <api-token>'
'https://api.tailscale.com/api/v2/tailnets?limit=100'
# Use the returned cursor value in subsequent requests:
curl -H 'Authorization: Bearer <api-token>'
'https://api.tailscale.com/api/v2/tailnets?limit=100&cursor=
<cursor>'
```

Page through all tailnets in a large organization by following the returned `cursor` until it is empty.

```
$ curl -H 'Authorization: Bearer <api_key>'
'https://api.tailscale.com/api/v2/tailnets?limit=100&cursor=
<cursor>'
```

— Names exact query params, default limit, response fields, and runnable curl example.

changelog-20260819-09cb7947 Tailnet list API pagination (2026-08-18)

└──▷ **BREAKING ON UPGRADE**

! The list tailnets endpoint now returns only the first 100 results by default for organizations with more than 100 tailnets; callers that previously received all tailnets in a single response must now paginate using the `cursor` parameter to retrieve the full set.

! The list tailnets endpoint now returns only the first 100 results by default for organizations with more than 100 tailnets; callers that previously received all tailnets in a single response must now page through results using the `cursor` parameter.

WAS THIS USEFUL? 👍 👎

◆ Third-Party & Vendor Risk (TPRM)

Conveyor

Feedback ⓘ 4 RELEASES · SEEN 2026-08-19 NOTES ↗

Conveyor is a tool that packages Java applications into native installers for Windows, macOS, and Linux with automatic updates.

> Conveyor extended its Trust Center with a new feedback API and an AI-agent MCP connector, gave ConveyorAI persistent context across Slack thread follow-ups, and previewed a rebuilt Word experience for questionnaire answering in Labs.

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

01 **Trust Center feedback endpoint and MCP tool** NEW 90

Adds a `GET /api/v2/trust_center/feedbacks` endpoint returning Trust Center survey responses including ratings, comments, and source, with filtering by source, rating, and date range plus standard pagination. The same data is exposed as an MCP tool so AI agents can read survey responses directly, with access to both the endpoint and MCP tool governed by the Trust Center Designer permission set.

.
.
.
.

— Exact endpoint, method, filters and permission set named

🔌 Trust Center Feedback Data Now in th…

02 **Trust Center AI Connector via MCP server** NEW 80

A new MCP server for Trust Centers lets visitor AI agents pull documents and ask questions without human involvement, enabled via the Enable visitors to connect via MCP toggle in Settings > Trust Center > Trust Center Agent. It enforces existing access groups, NDA gates, and gating rules against the authenticating visitor email even when access is agent-driven, logs every agent login, document access, and question the same as a human visit, and supports scheduled reassessments by the customer's agent.

.
.
.
.

Enable the MCP server so your customers' AI agents can programmatically run vendor risk reviews against your Trust Center.

> 📍 In the Conveyor console, go to Settings › Trust Center › Trust Center Agent and flip on Enable visitors to connect via MCP.

— Names exact toggle and settings path, mechanism fully described

🤖 Trust Center AI Connector: Let Your …

03 Editable Word question outline with instant export NEW

75

A Labs opt-in Word experience renders security questionnaire documents in-browser exactly as they appear in Word, builds an editable question outline, and lets users remove misdetected questions or drag to rescan sections. AI-drafted answers are placed back into the document in one to two minutes with blue highlights marking AI-touched content that clear on edit, and export is now instant, replacing the previous ten-to-fifteen minute wait.

— Clear mechanism and before/after timing, but Labs opt-in only

📄 A New Word Experience (Labs Only): E…

THINNER COVERAGE BELOW

04 Persistent Slack thread context for ConveyorAI IMPROVED

55

ConveyorAI now retains full conversation history within a Slack thread so follow-up questions carry the product line and prior answers forward without restarting context. It is enabled by default for every Slack integration and can be turned off in Organization Settings.

— Explains behavior and opt-out location, no other named surfaces

💬 Follow-Up Questions in Slack Threads…

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Promptfoo

Feedback ⓘ 1 RELEASE · SEEN 2026-08-19 NOTES ↗

Promptfoo is a testing framework for evaluating and benchmarking language model prompts and AI applications across multiple providers.

> Promptfoo added a new plugin to test AI systems for susceptibility to unverifiable claims.

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

01 **Unverifiable Claims Detection plugin** NEW 26

A new plugin tests AI systems' susceptibility to unverifiable claims.

— Names the plugin but no config key or usage detail. snapshot-20260819

WAS THIS USEFUL? 👍 👎

ToolHive

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

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.44.0 strengthens its outgoing authentication and delegation model, adding RFC 8693 actor_token support and closing SSRF gaps in OAuth flows.

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

01 actor_token support in RFC 8693 token exchange NEW

52

Outgoing authentication now supports the `actor_token` parameter in RFC 8693 token exchange flows, enabling richer delegation chains when ToolHive exchanges tokens on behalf of a user or service.

— Names the exact RFC and parameter but no usage steps v0.44.0

02 SSRF hardening for OAuth and DCR endpoints IMPROVED

35

OAuth token and DCR (Dynamic Client Registration) endpoints are now guarded against SSRF attacks, hardening the delegation/authentication security posture.

— States what's protected but not the mechanism v0.44.0

WAS THIS USEFUL? 👍 👎

SecureAgentics Adrian

Feedback ⓘ 1 RELEASE · 2026-08-18

CODE ↗

Open-source runtime AI agent security tool - monitors and controls AI agents, catching malicious tool use, prompt injection, and policy drift in real time, before the agent acts.

> Adrian's Python SDK v1.2.0 adds chain-of-thought capture to LLM event payloads, extending reasoning extraction across LangChain and TypeScript OpenAI SDK integrations.

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

01 **Reasoning capture in LLM event payloads** NEW 83

Adds a `reasoning` field to the LLM event payload emitted by `on_llm_end` in `sdk/python/adrian/handler.py`, captured via a new _extract_reasoning() helper. The TypeScript OpenAI SDK integration gains the same capability, collecting `delta.reasoning_content` from OpenAI-compatible streaming responses into a `reasoning` string alongside the existing `output`. Capture now spans multiple block layouts: `summary` lists (OpenAI Responses / LangChain default), bare `reasoning` strings (LangChain `output_version='v1'`), and Anthropic `thinking` strings.

— Names exact fields, files, helper function and supported block formats.

python-sdk-v1.2.0

WAS THIS USEFUL? 👍 👎

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-19 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

> OpenShell added OCSF AI inference event emission and bumped its OCSF schema to v1.8.0.

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

01 **OCSF AI inference event emission** NEW 70

OpenShell now emits AI inference events using the `ai_operation` profile on the OCSF `ApiActivity` class (event class 6003), with the OCSF schema bumped to v1.8.0.

— Names exact OCSF class, profile and schema version but no usage steps. v0.0.109

WAS THIS USEFUL? 👍 👎

NoLabs nono

Feedback ⓘ 1 RELEASE · 2026-08-19 NOTES ↗ CODE ↗

Sandbox any AI agent in seconds - zero setup, zero latency.

> nono v0.74.0 adds remote session connect and listing through nono-console, lets sandbox policies explicitly pass named environment variables through the dangerous-var blocklist, and makes supervised-mode approval prompts routable to configurable backends.

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

01 Remote session connect and listing via nono-console NEW

88

Adds `nono connect` to attach the current terminal to a session hosted by nono-console, accepting a session ID (e.g. `local:host:abc123`) or a URL directly. Extends `nono ps` with a `--remote` flag to list sessions hosted by the enrolled tenant's nono-console, combinable with `--all` and `--json`.

List all live and exited sessions running on the enrolled tenant's nono-console — useful for incident triage across a team.

```
$ nono ps --remote --all --json
```

Attach your terminal directly to a known remote session by its global session ID.

```
$ nono connect local:host:abc123
```

— Exact commands, flags and session-ID format given, fully runnable. v0.74.0

02 export_env pass-through for sandbox environment variables

NEW

85

Adds an `export_env` field to tool-sandbox `CommandPolicyConfig` and a top-level `session_export_env` list in `CommandPoliciesConfig`, letting callers explicitly pass named environment variables (exact names, trailing-`*` prefixes, or bare `*`) through to child processes, bypassing the dangerous-var blocklist.

— Names exact config keys and matching syntax, no usage example given. v0.74.0

THINNER COVERAGE BELOW

03 Configurable approval backends for supervised mode NEW

55

Adds a configurable approval backend for supervised-mode prompts via `approval_backends` in the profile, letting operators route capability-approval decisions to different backends.

— Names config key but omits backend types or setup steps. v0.74.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Syslifters SysReptor

Feedback ⓘ 1 RELEASE · 2026-08-19

NOTES ↗

A customizable and powerful penetration testing reporting platform for offensive security professionals. Simplify, customize, and automate your pentest reports with ease.

> SysReptor 2026.68 adds encrypted note file downloads, AI agent question-asking during interactions, default note structures for HTB and OffSec designs, and retest status color customization, while removing the deprecated Excalidraw plugin.

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

01 **Improved translation selection in scanimport plugin** IMPROVED

45

Improves template translation selection logic in the `scanimport` plugin.

— Names the plugin but not the exact logic change 2026.68

02 **Default note structures for HTB and OffSec designs** NEW

40

Adds default note structures to the HTB, OffSec, and demo designs, giving pentesters pre-built note templates for these engagement types.

— Names the three designs but not how to select them 2026.68

03 **Encrypted note file downloads** NEW

35

Note file downloads can now be sent through an encrypted channel to prevent intercepting proxies from inspecting the downloaded content.

— Names the capability but no config or flag to invoke it 2026.68

04 **Session cookie rotation after shared note auth** IMPROVED

35

Rotates the session cookie after shared note password authentication, reducing risk from session fixation on shared note access.

— Explains security rationale but no further mechanism detail 2026.68

05 **Excalidraw plugin removed** BREAKING — 35

The deprecated Excalidraw plugin has been removed.

— Names removed plugin but no migration guidance given 2026.68

06 **AI agent can ask user questions** NEW — 30

The AI agent can now ask the user questions during agent interactions, allowing interactive back-and-forth during automated tasks.

— States behaviour with no mechanism or example given 2026.68

07 **Customizable retest status colors** NEW — 30

Allows customizing retest status colors in the UI.

— Bare UI feature description, no path or options listed 2026.68

08 **Markdown preview rendered in iframes** IMPROVED — 25

Renders markdown preview content in iframes.

— One-line change with no explanation of impact 2026.68

└──▷ BREAKING ON UPGRADE

! The deprecated Excalidraw plugin has been removed.

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 19, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue