THE CYBER TOOLCHAIN NO. 063

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 3, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | **EVERYTHING**

Every tool in the issue, in full. all 126 changes

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 23 of 24 tools

▾ // THE BRIEF HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: **Opus 5** | GPT-5.5

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out. Fleet can now hand out managed local Windows admin accounts, so break-glass admin rights stop living in a

vim ? help

shared password. JFrog's Zero-Touch Remediation swaps a vulnerable package for a compliant version instead of filing a ticket. Falcon Onum runs detection in-stream before data lands, so noise never gets billed or stored.

PROTECT

Stop handing every Windows user a permanent local admin

Fleet 4.91.0 manages local Windows admin accounts and pins macOS/iOS/iPadOS updates to `latest`, which removes the spreadsheet of shared admin passwords and the quarterly re-pin of version numbers as new OS releases land. Separately, OpenAI's Trusted Access for Cyber program now mandates hardware-backed passkeys and disables weaker methods outright — phishable fallbacks are the usual reason an MFA rollout doesn't reduce account takeover.

Fleet · Yubico

RESPOND

Let the fix ship without a human writing the patch PR

JFrog's Zero-Touch Remediation automatically swaps vulnerable packages for compliant versions, collapsing the usual find-triage-bump-test loop that leaves known-bad dependencies in place for weeks. ZeroPath's AppSec Assistant takes the opposite tack on trust: it proposes a reviewable plan before acting, so an AI-driven change is approved rather than discovered after the fact.

JFrog · ZeroPath

DETECT

Drop the noise before it becomes ingest cost

Falcon Onum's Certified Data Pipelines filter and run detection in-stream before events land in storage, which moves the economics of what you can afford to collect. Tenzir v6.15.0 adds Splunk HEC acknowledgements so a pipeline knows delivery actually succeeded rather than assuming it, plus up to 9x faster event branching.

CrowdStrike Falcon Onum · Tenzir

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

T3MP3ST

Feedback ⓘ changes since 2026-08-24 CODE ↗

autonomous red teaming platform; multi-agent offensive-security meta-harness

T3MP3ST adds social OSINT tools, headless-browser XSS probing, WAF-aware scanning, and Windows parity

Adds `username_search` tool to the `social_osint` family: checks a username's existence across GitHub, GitLab, Reddit, Hacker News, Wikipedia, Steam, Keybase, VK, and Telegram using unauthenticated public surfaces — no API keys required (Sherlock-style).

Adds `telegram_lookup` tool: resolves a @username via the Telegram Bot API `getChat` call, returning public info (name, description, type, member count); reads the `TELEGRAM_BOT_TOKEN` environment variable when present and degrades gracefully without it.

Adds `email_format` tool: generates corporate email address candidates from a person's name and validates the domain's MX records via live DNS probe — no email sending.

Adds a headless-Chromium browser probe (`browser.ts`) that captures rendered DOM, page title, cookie flags (never values), and console errors, and performs a real XSS check by injecting a payload into a query parameter and detecting both reflection and alert-dialog execution.

Adds `resolveSystemPrompt` export to the `prompts/index` module, enabling dynamic system-prompt resolution for operators.

WAS THIS USEFUL? 👍 👎

Strix

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

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

> Strix v1.6.1 refreshes its recommended LLM model list and hardens MCP integrations against transient call failures.

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

01 **Expanded recommended model list** NEW 70

Strix adds `openrouter/z-ai/glm-5.3` as the new top recommended model, alongside GLM-5.3, Fable 5.1, and Gemini 3.7 Flash added to the recommended model list. Users can switch models by setting the `STRIX_LLM` environment variable, e.g. `export STRIX_LLM="openrouter/z-ai/glm-5.3"` alongside `LLM_API_KEY`, then running `strix --target ./app-directory`.

Switch to the new top-pick model for a pentest scan without editing config files.

```
$ export STRIX_LLM="openrouter/z-ai/glm-5.3"
export LLM_API_KEY="your-api-key"
strix --target ./app-directory
```

— Names exact model IDs and a runnable command v1.6.1

THINNER COVERAGE BELOW

02 **MCP call errors no longer quarantine connection** IMPROVED 30

MCP call errors no longer quarantine the connection, keeping MCP integrations live after transient failures.

— States behaviour change but no reproduction steps v1.6.1

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

evil-winrm

Feedback ⓘ 1 RELEASE · 2026-09-03 NOTES ↗ RANK

The ultimate WinRM shell for hacking/pentesting

> evil-winrm v4.1 adds explicit shell exit options, a keepalive to stop idle WinRM disconnects, and updated Ruby compatibility.

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

01 **Ruby 3.3 and 4.0 compatibility** IMPROVED 55

Adds `readline-ext` as a required dependency to restore remote path completion compatibility with Ruby 3.3 and newer, and adds Ruby 4.0 support.

— Names dependency and versions but no usage steps v4.1

02 **Explicit shell exit commands** NEW 50

Adds a `quit` command inside the shell as an explicit exit mechanism, and adds Ctrl+D support as an alternative way to exit the shell session.

— Names exact command and key combo, no further mechanism given v4.1

03 **Keepalive for idle WinRM shells** NEW 35

Adds keepalive to prevent Windows from disconnecting idle WinRM shells.

— States purpose but no config detail or mechanism v4.1

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

ZeroPath

Feedback ⓘ 1 RELEASE · SEEN 2026-09-03 NOTES ↗ RANK

AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives. Automated vulnerability detection, instant fixes, and seamless CI/CD integration for modern DevOps teams.

> ZeroPath's biggest shift is an AI AppSec Assistant that now proposes reviewable plans before taking action, backed by a hosted remote MCP server, continuous default-branch scanning, and a broad access-sync overhaul spanning GitHub and Bitbucket. The platform also picked up Gerrit and CVS repository support, four new compliance reports (SOC 2, SLA Compliance, Organization Summary, Monthly Security), and CVSS v4.0 scoring.

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

01 Bitbucket and GitHub access sync with permission mapping

90

BREAKING

A Bitbucket Access Sync section in team settings imports Bitbucket Cloud user groups and Bitbucket Server groups as ZeroPath teams, with optional hourly auto-sync and a Re-sync now / Sync now button for one-off imports; a permission mapping editor warns when a custom mapping grants fewer permissions to a higher VCS role than a lower one, and unmatched Bitbucket group members appear in the same manual-linking panel used for GitHub, tagged with a platform badge. This extends to a broader access-sync system importing teams from both GitHub and Bitbucket, mapping GitHub's five roles and Bitbucket's three onto thirty-nine permissions across ten groups, and restricting member visibility to only repositories they hold a grant on. Connecting an integration now requires the Create Integrations or Create Installations permission, and adding a repository by URL requires the Create Repositories permission.

Enable hourly Bitbucket group sync to automatically keep ZeroPath team membership in step with Bitbucket Cloud or Server groups.

📍 In ZeroPath, go to Team Settings › Bitbucket Access Sync and toggle on Enable auto-sync. A first sync is queued immediately; subsequent syncs run hourly.

Customize how Bitbucket roles map to ZeroPath permissions — for example, restricting Write-role users from triggering scans.

📍 In ZeroPath, go to Team Settings › Bitbucket Access Sync, click the permission mapping editor button, then adjust the permission matrix (VCS roles as columns,

ZeroPath permission groups as rows) and save. A warning is shown if a higher role ends up with fewer permissions than a lower one.

— Exact permission counts and button labels; UI-driven only.

Remote MCP server, Microsoft Teams, Ger…

02 **Remote MCP server with OAuth 2.1** NEW 85

A hosted remote MCP endpoint at `https://zeropath.com/mcp` uses OAuth 2.1 over the Streamable HTTP transport, letting clients like Claude, Cursor, and Codex connect via browser authentication with no install or API key; connections are read-only until changes are approved. A local `stdio` transport mode running the open-source Python server with API-token credentials remains available for scripted or self-managed setups.

Connect an AI coding assistant to your ZeroPath findings without installing anything or managing API keys.

```
$ # In Cursor, Claude, or any Streamable HTTP MCP client, set the MCP server URL to:
https://zeropath.com/mcp
# Then authenticate via the browser prompt that appears — no API key needed.
```

Point a Cursor or Claude MCP client at the hosted ZeroPath endpoint to start using it with no local install.

```
$ https://zeropath.com/mcp
```

— Exact endpoint and auth flow given; stdio config left generic.

Remote MCP server, Microsoft Teams, Ger…

03 **AI AppSec Assistant with Planning Mode and Skills** NEW 83

A per-repository chat-based AppSec Assistant answers security questions and executes automated actions on findings; before complex or irreversible actions (opening pull requests, bulk-updating findings) it proposes a reviewable plan listing Steps, Risks, and scope that must be approved, forceable via a Plan first toggle on any message and savable as a schedule, trigger, or playbook. Plans surface Scheduling and Trigger Suggestions (a cron expression, event trigger, or playbook activation) when a request is recurring or reactive. Organization-specific guidance is authored via the Skills tab (Early Access), with four built-in forkable skills for triage, dependency reachability, report writing, and patch review, and active skills are listed in the plan's `skills used` field.

Review and approve an AI-proposed plan before the assistant bulk-opens fix PRs for critical findings, so you can verify scope and risks before any irreversible action is taken.

📍 In the AppSec Assistant chat, type your request (e.g. 'Open fix PRs for all critical findings in this repo'). The assistant displays a plan with Steps, Risks, and scope. Review each field, then click Approve to proceed — or tick `Plan first` on the message before sending to always force a plan.

— Names toggle and skills tab, but no CLI or API surface.

`Remote MCP server, Microsoft Teams, Ger…`

04 **Patch submission status tracking** NEW 83

Fix PRs for approved patches show progress through `Queued`, `Submitting PR`, `Retrying`, `Submission failed`, `Cancelled`, and `PR opened` states, with specific failure reasons (e.g. missing push credentials) and a `Retry` button for unrecoverable failures; up to five automatic submission attempts are made before a PR is marked stopped, and `Cancel` / `Retry` actions appear directly on the issue.

Track the progress of an automated fix PR from submission through to merge, and manually retry if submission fails due to a resolvable issue such as missing push credentials.

📍 Open the finding panel after promoting a patch. Watch the status chip cycle through `Queued` → `Submitting PR` → `PR opened`. If it lands on `Submission failed`, read the displayed failure reason (e.g. `Missing push credentials`), reconnect the repository through your VCS integration, then click the `Retry` button.

— Named states and retry limit; UI-only actionability.

`Remote MCP server, Microsoft Teams, Ger…`

05 **Continuous Scanning on default branch** NEW 83

When enabled, ZeroPath runs a full scan on every meaningful commit to a repository's default branch (including merged pull requests) and incrementally re-examines only changed code on follow-up scans, automatically closing findings resolved by those commits; configurable per organization, tag, or repository in scanner settings, and available on GitHub, GitLab, Bitbucket Cloud, and Bitbucket Server.

Enable continuous scanning on a GitHub organization so findings automatically track the default branch between scheduled scans.

📍 1. In ZeroPath, go to Scanner Settings. 2. Select the organization, tag, or repository you want to configure. 3. Enable `Continuous Scanning` — subsequent merged pull requests and commits to the default branch will each trigger a full scan automatically.

— Named scope levels and VCS support list; toggle in Scanner Settings.

Remote MCP server, Microsoft Teams, Ger…

06 **New compliance and audit reports** NEW 80

Four organization-level report cards are added, visible only to users with organization-level permission: Monthly Security Report (DOCX, with an AI-generated executive summary plus per-repository scan activity, severity breakdown, and top vulnerability families), Organization Summary Report (PDF), SLA Compliance Report (PDF, tracking which findings are within or past configured remediation windows), and SOC 2 (PDF, compliance-focused metrics and scan activity structured for audit evidence). All reports carry a Scope field on the cover page or header showing either All repositories (N) or the specific `tags` / `repositories` selected at generation time; by default a report covers the entire organization.

Generate a SOC 2 audit report scoped to a specific team tag so only the relevant repositories appear in the cover-page Scope field.

📍 In the Reports tab, open the SOC 2 report card, select the desired tag(s) or repository names under the scope picker, then click Generate. The produced PDF cover page will read the specific tags and repository names you chose rather than All repositories.

— Names each report format and scope mechanism; no API endpoint given.

Remote MCP server, Microsoft Teams, Ger…

07 **Dashboard charts and reporting UI** NEW 80

A new Issues By PR Author chart groups open issues from each repository's latest merged pull request scan by author, excluding unmerged PR scans and full-scan-only findings. A new Exposure Window chart shows the average days open issues of each severity were exposed per period, with severity lines breaking rather than dropping to zero when none are open. The date range picker adds a Last 1 year preset and Quarterly granularity alongside daily/weekly/monthly options. The Top Issues widget and scored-findings tables sort unscored findings to the bottom and support up to 500 findings with severity, score, title, repository, vulnerability class, status, and detection-date columns; clicking any row opens a side pane with full issue details, and actions like marking false positive refresh the report in the background without losing filters, scroll position, or selection.

Review which PR authors introduced the most open issues shipped to production, filtering out unmerged PR noise.

📍 In the Reports tab, open the Issues By PR Author chart. The chart automatically shows open issues from each repository's latest merged pull request scan only — unmerged PRs and full-scan findings are excluded — grouped by author.

— Names both new charts and exact table limit; UI-only.

Remote MCP server, Microsoft Teams, Ger…

08 **Jira and Linear ticket integration** BREAKING 69

A distinct Issue Tracker Export permission group in team settings covers Create Jira tickets from findings and Create Linear tickets from findings, separate from Issue Management. Jira and Linear ticket links are now included in CSV exports and the export endpoint, and an on-demand Jira ticket sync preview shows what a linked ticket will do before it's applied. Creating Jira and Linear tickets is now a per-repository permission that must be explicitly granted or ticket creation is blocked.

— Named permission group and export surfaces; endpoint path unnamed.

Remote MCP server, Microsoft Teams, Ger…

09 **Microsoft Teams integration** NEW 67

Adds Microsoft Teams integration: paste an incoming webhook URL into ZeroPath to receive scan notifications and vulnerability alerts as Adaptive Cards in Teams channels, with no Azure AD app registration or tenant administrator required.

— Exact setup step (paste webhook URL) with no prerequisites.

Remote MCP server, Microsoft Teams, Ger…

10 **Feedback disposition changes** BREAKING 62

The feedback list gains filter chips to narrow by disposition type; the Pending review disposition label is renamed to Needs review, and the Non-exploitable disposition now immediately removes a finding from the open issues queue (same as False Positive or Accepted Risk), whether set by a human or an automated agent.

— Names exact renamed label and queue behaviour change. product docs

11 **Gerrit repository support** NEW 62

ZeroPath adds Gerrit change scanning via poll-based triggering (no webhook required): it imports Gerrit projects for full scans, scans each open change as new patch sets arrive, and posts findings as inline review comments on diff lines / as a review message on the patch set.

— Names polling mechanism and comment placement; no config keys given.

Remote MCP server, Microsoft Teams, Ger…

12 **Dependency and SCA findings surfacing** IMPROVED 62

A new dependency findings tab on the Issues page supports search by CVE, advisory, and package; completed pull request scans link to their findings list with a separate dependency link, and the check run's Details link now points directly to the

dependency view when a scan finds only SCA findings, avoiding a dashboard view that appears empty.

— Names search fields and link behaviour; no endpoint given.

Remote MCP server, Microsoft Teams, Ger…

13 **Stale PR scan auto-cancellation** IMPROVED 60

ZeroPath automatically cancels in-progress PR scans for older commits when a newer commit is pushed to the same pull request, across GitHub, GitHub Enterprise Server, GitLab, Bitbucket, Azure DevOps Services, and Gerrit.

— Names all six supported VCS platforms; behaviour is automatic. product docs

THINNER COVERAGE BELOW

14 **CVSS v4.0 vector support** NEW 58

Adds CVSS v4.0 vector support for assessing subsequent-system impact, with severity labels following CVSS rating bands: Critical ≥ 9.0, High ≥ 7.0, Medium ≥ 4.0, Low ≥ 0.1.

— Exact severity thresholds given; no UI or config path.

Remote MCP server, Microsoft Teams, Ger…

15 **CVS repository support** NEW 52

Adds CVS repository support: import modules by server path over SSH or pserver, enabling full scans, scheduled scans, dependency analysis, and secrets scanning.

— Names transport protocols; no further config detail.

Remote MCP server, Microsoft Teams, Ger…

16 **Repository-size counting for more languages** NEW 45

Adds repository-size counting for Terraform, Svelte, Vue, Solidity, Elixir, Dart, protobuf, and Genero, ensuring those repositories receive analysis depth matched to their size.

— Lists all eight languages; no config or command given.

Remote MCP server, Microsoft Teams, Ger…

17 **Duplicate repository detection** NEW 45

Adds duplicate detection when adding repositories, surfacing likely duplicates from host migrations with a git history probe for each match.

— Names detection mechanism briefly; no UI path given.

Remote MCP server, Microsoft Teams, Ger…

18 **Expanded final validation coverage** IMPROVED 43

Extends final validation to cover up to 1,000 findings per scan, twenty times more than before.

— Exact numeric limit given; no mechanism explained.

Remote MCP server, Microsoft Teams, Ger…

19 **Secrets scanning coverage reporting** IMPROVED 43

Secrets scanning now reports the coverage achieved and names any file too large to read within the scan's time budget.

— Describes new report content; no example output given.

Remote MCP server, Microsoft Teams, Ger…

20 **Supply Chain scope control** NEW 40

Adds a Supply Chain scope control that narrows all tabs to a selected set of repositories, teams, and tags.

— Names scope dimensions; no exact UI path given.

Remote MCP server, Microsoft Teams, Ger…

21 **"What's new" release panel** NEW 34

Adds a What's new panel in the top bar showing the three most recent releases with unread indicators.

— Bare UI description; exact location given. Remote MCP server, Microsoft Teams, Ger…

22 **GitHub Enterprise Cloud with data residency** NEW 32

Adds GitHub Enterprise Cloud support with data residency, connecting on its own host.

— Bare mention, no configuration steps given. Remote MCP server, Microsoft Teams, Ger…

└──▷ BREAKING ON UPGRADE

! The Pending review disposition label is renamed to Needs review in the feedback list filter chips.

! Creating Jira and Linear tickets is now a per-repository permission; organizations using access control must explicitly grant it or ticket creation will be blocked.

! Connecting an integration now requires the Create Integrations or Create Installations permission; adding a repository by URL requires the Create Repositories permission.

WAS THIS USEFUL? 👍 👎

Codacy

Feedback ⓘ 1 RELEASE · SEEN 2026-09-03 NOTES ↗ RANK

Documentation homepage for the Codacy automated code review tool.

> Codacy added Oxlint, a Rust-based JavaScript/TypeScript linter, as a fully integrated analysis engine with broad ESLint-compatible rule coverage and automatic repository configuration detection.

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

01 **Oxlint added as supported analysis engine** NEW 80

Codacy adds Oxlint as a supported engine for JavaScript and TypeScript, covering correctness, style, and best-practice rule categories including ESLint core rules, `typescript-eslint`, React ecosystem (`react`, `react-hooks`, `react-refresh`, `react-perf`), `@next/eslint-plugin-next`, `eslint-plugin-unicorn`, testing frameworks (`jest`, `vitest`), `eslint-plugin-vue`, `import`, `jsdoc`, `jsx-a11y`, `n` (Node), and `promise`. It supports repository-level configuration via `.oxlintrc.json`, `.oxlintrc.jsonc`, `oxlint.config.json`, or `oxlint.config.jsonc` files, which Codacy detects and applies automatically during analysis.

— Names engine, rule sets, and exact config file names supported.

Adding Oxlint as new supported tool Sep…

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ 1 RELEASE · SEEN 2026-09-03 NOTES ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep expanded its public API with partial-scan visibility and new Agentic Workflows v2 endpoints, while adding policy automation, developer triage controls, and read-only token permissions alongside smaller registry and workflow UI improvements.

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

01 **Read-only Agent (CI) API token permission** NEW 65

Adds a read-only permission option when creating an Agent (CI) Web API token, allowing data reads without write access.

Provision a read-only CI token so a reporting pipeline can pull findings data without any risk of modifying policies or triage state.

📍 In Semgrep AppSec Platform, go to Settings › Tokens › Create Token, set type to `Agent (CI)`, and select the `Read-only` permission before saving.

— Exact navigation steps provided to create the token. snapshot-20260903

02 **Require exception requests policy setting** NEW 60

Adds a `Require exception requests` setting to policies, letting teams set developer triage permissions between no triage and full triage and require approval before a blocking finding can be ignored.

— Explains triage levels but not exact UI path or config key. snapshot-20260903

THINNER COVERAGE BELOW

03 **is_partial_scan field in public API** NEW 55

Adds `is_partial_scan` field to the public API response and as a request filter, indicating whether a scan ran only part of the rule set.

— Names the field but not the endpoint path or full schema. snapshot-20260903

04 **Incident scan type filter on Scans page** NEW 55

Adds `Incident` as a scan type option in the `Details > Scans` page scan type filter, alongside `Full` and `Diff`, for incident-scoped rule checks.

— Exact UI location given but no further mechanism. snapshot-20260903

05 **Registry confirmation dialog for policy additions** IMPROVED 55

Semgrep Registry now shows a confirmation dialog before adding rules to a detection policy, displaying rule count, target policy, and an immediate-effect notice.

— Describes dialog contents but not an exact navigation path. snapshot-20260903

06 **Search workflow issues by title and description** IMPROVED

45

Workflow issues can now be searched by title and description, including partial matches, via the magnifier filter bar.

— Names the filter UI element but little mechanism detail. snapshot-20260903

07 **Multi-webhook remediation policy automations** IMPROVED

40

Supports multiple webhook actions on a single remediation policy automation, enabling richer automated response workflows.

— Describes capability but no config keys or setup steps. snapshot-20260903

08 **Agentic Workflows endpoints in v2 API** NEW

35

Adds Agentic Workflows endpoints to the public v2 API reference.

— No endpoint paths, methods, or schema given. snapshot-20260903

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ 2 RELEASES · 2026-09-02 BLOG ↗ RANK

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

> GitHub Copilot administration gains two enterprise controls this window: configurable default AI models per team and enforcement of content exclusion policies in the Copilot app and CLI.

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

01 Per-team default model in enterprise-managed settings NEW

85

Enterprise-managed settings add a `model` key, configurable per team in `team-mappings.json`, to set a default Copilot model for each enterprise team; the key is marked `overridable` so teams can choose their own preferred model. A global enterprise default model applies to all Copilot users, with non-mapped users inheriting that default, and the feature is now generally available across Copilot Business and Copilot Enterprise in the GitHub Copilot app, Copilot CLI, and Visual Studio Code.

— Names exact config key, file, and scope Enterprise-managed settings support any…

THINNER COVERAGE BELOW

02 Content exclusions enforced in Copilot app and CLI IMPROVED

43

The Copilot app and Copilot CLI now respect content exclusion policies set by enterprise, organization, and repository administrators, preventing excluded files from being used as context in agentic workflows.

— Explains behavior but no config surface named Content exclusions generally available …

WAS THIS USEFUL? 👍 👎

noir

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

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

> Noir v1.3.1 focuses on diagnostic quality: plain-text output now matches structured formats, route findings gain precise source line numbers across more frameworks, and large OpenAPI documents no longer get silently dropped.

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

01 Line numbers for route declarations across frameworks IMPROVED

65

Noir now records source line numbers for route declarations so findings point to the exact line, covering Rails `config/routes.rb`, OpenAPI/Swagger operations, Symfony `#[Route]` attributes and `routes.yaml`, Vert.x, and Spring reactive `RouterFunction` routes.

— Names every framework covered but no invocation example v1.3.1

02 Larger parsing budget for big OpenAPI documents IMPROVED

60

OpenAPI documents over 10 MB now get their own 64 MB parsing budget instead of being silently dropped, as they previously were.

— Clear before/after with exact size thresholds v1.3.1

THINNER COVERAGE BELOW

03 Protocol and internal badges in plain text output IMPROVED

52

The plain text report now shows the protocol (`kafka` , `mqtt` , `amqp` , ...) and an `[internal]` badge, matching what structured formats like `-f json` and `-f yaml` already displayed.

— Names formats and flags but plain output not runnable directly v1.3.1

WAS THIS USEFUL? 👍 👎

JFrog

Feedback ⓘ 2 RELEASES · 2026-09-02 BLOG ↗ RANK

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

> JFrog's swampUP 2026 announcements extend platform governance into the agentic software supply chain, headlined by Zero-Touch Remediation that automatically swaps vulnerable packages for compliant versions, new governed registries for agent packages, plugins and skills, an AI-powered policy-as-code playground, and prompt-to-release traceability for compliance.

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

01 Zero-Touch Remediation via Compliant Version Replacement NEW

72

JFrog Zero-Touch Remediation automatically evaluates all available fixes for a vulnerability in a customer's environment and selects the best one based on policy, environment, and risk tolerance. It works through Compliant Version Replacement, swapping a vulnerable package version for a patched one mid-pipeline by evaluating multiple upstream sources (Chainguard , Maven Central , and others) against the customer's own policy, with JFrog Artifactory as the system of record.

— Explains mechanism and named sources but no exact API or config.

Extending the Single Source of Truth to… | Live From the Show Floor: swampUP 2026 …

02 Governed registries and resolution for agent packages and plugins

NEW

68

New APM Repository provides a governed registry for agent package installs, replacing open-internet pulls via `apm install`. New Plugins Registry delivers a scanned, auditable source for first- and third-party coding agent plugins, replacing default public marketplaces. New Agent Package Resolution provides a trusted, policy-enforced path for agents to resolve dependencies through the JFrog Platform.

— Names the `apm install` command and each surface's role.

Extending the Single Source of Truth to…

03 JFrog Agent Plugin for coding agents NEW

63

New JFrog Agent Plugin integrates coding agents — Claude Code , Cursor , VS Code Copilot , Kiro , Devin , OpenCode , and Codex — natively with the JFrog Platform.

— Names seven specific agent tools integrated natively.

Extending the Single Source of Truth to…

04 **AI-powered policy-as-code playground** NEW 60

Adds an AI-Powered Policy as Code Playground that converts plain-language intent into enforceable, versioned, Rego-based policy automatically, requiring no manual policy authoring.

— Names Rego and describes the conversion mechanism.

Live From the Show Floor: swampUP 2026 …

THINNER COVERAGE BELOW

05 **Traffic Controller integrations with network proxies** NEW 59

New Traffic Controller integrations with Netskope, Zscaler, and Cloudflare redirect public-repo package requests through JFrog at the network layer.

— Names three vendors and the network-layer mechanism.

Extending the Single Source of Truth to…

06 **Agentic remediation via AutoPR pull requests** NEW 50

AutoPR with Smart Remediation automatically pushes verified, context-aware fix guidance directly to pull requests. For first-party code this is Agentic Remediation, fixing vulnerabilities at the source level with a developer reviewing and approving changes before merge — distinct from Zero-Touch Remediation's fully automated mid-pipeline swap.

— Names the mechanism but no command or config detail.

Extending the Single Source of Truth to… Live From the Show Floor: swampUP 2026 …

07 **Repo-level policies without project setup** NEW 49

Adds Repo-Level Policies that enforce evidence-based governance directly on repositories without requiring project, application, or lifecycle-stage setup.

— Describes the before/after setup requirement clearly.

Live From the Show Floor: swampUP 2026 …

08 **Post-release governance for compliance windows** NEW 45

Adds Post-Release Governance, extending compliance monitoring and vulnerability remediation tracking through the support window after a release ships to address European regulatory requirements.

— Names the regulatory driver but no specific mechanism.

Live From the Show Floor: swampUP 2026 …

09 **Skills Scanning for agent skill files** NEW 43

New Skills Scanning uses semantic security scanning to detect malicious behavior in skill files and block them before they reach a developer's machine.

— States the mechanism briefly but no config or scope detail.

Extending the Single Source of Truth to…

10 **Agent Guard AI asset enforcement** NEW 43

New Agent Guard (AI Asset Enforcement) ensures only AI assets approved in AI Catalog are consumed by agents across coding environments.

— Names AI Catalog link but no enforcement mechanism detail.

Extending the Single Source of Truth to…

11 **AI Catalog expanded to skills, plugins and Agent Packages**

IMPROVED 41

Expanded JFrog AI Catalog now governs skills, plugins, and Agent Packages (APM) in addition to previously supported AI assets.

— States scope expansion but no mechanism or config detail.

Extending the Single Source of Truth to…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog

Feedback ⓘ

1 RELEASE · 2026-09-02 | NOTES ↗ | CODE ↗ | RANK

trufflehog discovers, classifies, and validates leaked credentials

> TruffleHog v3.97.2 adds a new AWS Secrets Manager source, a Weights & Biases V2 detector, GitHub token-type differentiation for faster remediation, Jenkins resource enumeration, and more resilient git cloning against GitHub rate limits.

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

01 **GitHub token-type differentiation in detector** IMPROVED — 80

Adds `token_type` and `remediation` fields to `ExtraData` in the GitHub V2 detector, differentiating between PAT (`ghp_` , `github_pat_`), OAuth (`gho_`), and GitHub App (`ghu_` , `ghs_` , `ghr_`) token types so incident responders are directed to the correct revocation path.

— Names exact fields and token prefixes plus behavioral benefit. v3.97.2

02 **Weights & Biases V2 detector** NEW — 65

Adds a V2 detector for Weights & Biases, enabled via the `WeightsAndBiasesV2DetectorEnabled` feature flag, expanding credential coverage for the ML platform.

— Names the feature flag but no usage detail beyond enabling it. v3.97.2

03 **AWS Secrets Manager source** NEW — 60

Adds the `SOURCE_TYPE_AWS_SECRETS_MANAGER` source type, enabling TruffleHog to scan AWS Secrets Manager for leaked credentials.

— Names the source type but no scan config or CLI detail. v3.97.2

THINNER COVERAGE BELOW

04 **Automatic git clone retry on rate limits** IMPROVED — 55

Retries `git clone` automatically on secondary rate-limit responses (bare 403/429), reducing scan failures against large GitHub targets.

— Explains trigger and benefit but no configurable retry options given. v3.97.2

05 **Jenkins source enumeration** IMPROVED — 50

Implements `Enumerate` for the Jenkins source, allowing TruffleHog to discover and iterate Jenkins resources during scans.

— Names the mechanism but not what resources are discovered. v3.97.2

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ 1 RELEASE · 2026-09-03 BLOG ↗ RANK

Try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

> Elcomsoft Quick Triage 2.2 ships a cross-artifact timeline, a new plugin architecture for artifact parsing, and a file-system-metadata snapshot capability, alongside extended MSA password attacks, full-text search and BitLocker key extraction.

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

01 **File system snapshot artifact** NEW 45

A new file system snapshot capability copies file system metadata without extracting file contents, and is available as a new artifact type for collection and analysis.

– Clear mechanism (metadata-only copy) but no invocation steps.

Elcomsoft Quick Triage 2.2 adds a timel…

02 **Cross-artifact timeline view** NEW 40

A new Timeline view merges events from every timestamped artifact into a single chronology, letting investigators follow activity across different artifact types in one merged event view.

– Describes purpose but no UI path or mechanism detail.

Elcomsoft Quick Triage 2.2 adds a timel…

03 **Plugin architecture for artifact parsers** NEW 40

A new plugin architecture rebuilds artifact support in Quick Triage, enabling third-party and custom artifact parsers and extensible artifact collection.

– Explains extensibility purpose, no API or plugin format given.

Elcomsoft Quick Triage 2.2 adds a timel…

04 **MSA password attacks** NEW 40

Extended password recovery capabilities now include attacks against Microsoft Account (MSA) passwords.

– Names the account type targeted but no attack method detail.

Elcomsoft Quick Triage 2.2 adds a timel…

05 **Full-text search format and archive support** IMPROVED 40

Full-text search now parses OpenDocument format files and supports recursive archive parsing, extending overall search capabilities.

— Names formats supported but no scope limits or usage steps.

Elcomsoft Quick Triage 2.2 adds a timel…

06 **Extended BitLocker key extraction** IMPROVED 20

Extended BitLocker key extraction capabilities in Quick Triage 2.2.

— Only a bare mention with no detail on what changed.

Elcomsoft Quick Triage 2.2 adds a timel…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ IDS / IPS

CrowdSec

Feedback ⓘ 1 RELEASE · 2026-09-03 NOTES ↗ RANK

CrowdSec - the open-source and participative security solution offering crowdsourced protection against malicious IPs and access to the most advanced real-world CTI.

> CrowdSec's v1.8.1 release adds a small helper function for bot detection scenarios that validates challenge cookies.

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

01 **HasValidChallengeCookie() helper for bot detection** NEW 50

Adds a HasValidChallengeCookie() helper function for use in bot detection scenarios, enabling challenge cookie validation checks.

— Names the exact function but no usage context or scenario example. v1.8.1

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Fleet

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

Open device management

> Fleet 4.91.0 ships managed local Windows admin accounts, latest OS version pinning for macOS/iOS/iPadOS update policies, per-fleet activity webhooks, an end-user-auth bypass flag for fleetd installers, direct-from-storage software installer downloads, Adobe Creative Cloud plugin inventory, and 29 new iOS/iPadOS device vitals, alongside a batch of smaller API, activity-log, and UI additions.

└──▷ WHAT SHIPPED · 20 FEATURES most completely described first ? what's the number?

01 **Bypass end-user auth for fleetd enrollment** NEW 95

Adds `--bypass-end-user-auth` flag to `fleetctl package` to generate fleetd installers that skip end-user authentication prompts on Linux and Windows hosts (requires fleetd v1.60.0 or higher). Adds `mdm.allow_orbit_end_user_auth_bypass` server setting, enabled by default, to control whether hosts that skip end-user authentication can enroll into a fleet that requires it; set to `false` to strictly enforce authentication for all Orbit enrollments.

Enforce end-user authentication for all Orbit enrollments, preventing hosts that skip MFA from joining a fleet that requires it.

yaml

```
mdm:
  allow_orbit_end_user_auth_bypass: false
```

Build a fleetd installer for a kiosk deployment where end users have already authenticated via another MDM and should not see another auth prompt.

```
$ fleetctl package --type=msi --fleet-
url=https://fleet.example.com --enroll-secret=<secret> --bypass-
end-user-auth
```

— Flag, setting, and runnable commands all given fleet-v4.91.0

02 **Per-fleet software inventory toggle** NEW 90

Adds `PATCH /api/v1/fleet/fleets/{id}` support for `{"features": {"enable_software_inventory": <bool>}}` to enable or disable software inventory collection per fleet; omitting the key leaves the stored value unchanged.

Disable software inventory for a specific fleet to reduce database load on a fleet used only for script execution.

```
$ curl -X PATCH https://fleet.example.com/api/v1/fleet/fleets/7
\
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"features": {"enable_software_inventory": false}}'
```

– Exact endpoint, payload, and curl example given `fleet-v4.91.0`

---

03 **Custom host vital variables in Android profiles** NEW

80

Adds support for `$FLEET_HOST_VITAL_<id>` custom host vital variables in Android configuration profiles, managed app configurations, and host name templates, with per-host expansion and automatic re-delivery when a vital value changes.

– Variable syntax and behaviour named, no example `fleet-v4.91.0`

---

04 **'Latest' OS version pinning for update policies** NEW

75

Adds support for setting `minimum_version` to `latest` with `deadline_days` in macOS/iOS/iPadOS OS update policy to automatically track the newest Apple-published version per device (Fleet Premium).

– Named fields and behaviour, no full example call `fleet-v4.91.0`

---

05 **Direct-from-storage software installer downloads** NEW

70

Adds `s3_software_installers_signed_url` configuration option to serve software installers, in-house apps, and bootstrap packages via GCS presigned URLs so clients download directly from object storage instead of routing through the Fleet server.

– Config key named; behaviour explained, no runnable example `fleet-v4.91.0`

---

06 **Apple Business Manager error visibility and device release** NEW

70

Adds `token_invalid` field to ABM token API responses and `dep_device_error` (human-readable message) to `GET /hosts/:id/dep_assignment` to surface why an Apple Business Manager device lookup or ABM token is failing. Also adds support for releasing devices from Apple Business Manager inside Fleet.

– Named endpoint and fields, no example call shown `fleet-v4.91.0`

---

07 **29 new iOS/iPadOS device vitals** NEW — 70

Adds 29 new iOS/iPadOS device vitals (battery level, accessibility settings, cellular technology, cloud backup status, MDM options, device attestation, and more) collected via the existing `DeviceInformation` MDM refetch and exposed in the host API response and a new View all vitals modal; BYOD enrollments are excluded.

— Mechanism and scope detailed, no field list or example `fleet-v4.91.0`

08 **Host identifiers on MDM enrollment activity** IMPROVED — 65

Adds `host_id` and `host_serial` fields to the `mdm_enrolled` activity for Apple (macOS, iOS, iPadOS) enrollments, and surfaces the activity on the host's activity timeline.

— Named fields and surface, no example `fleet-v4.91.0`

09 **Built-in variables in software scripts** NEW — 65

Adds support for Fleet built-in variables (e.g. `$FLEET_VAR_HOST_END_USER_IDP_USERNAME`) in software install, post-install, and uninstall scripts, resolved per host at execution time (Fleet Premium).

— Named variable and scope, no example usage `fleet-v4.91.0`

10 **Default fleet assignment for Windows enrollments** NEW — 65

Adds a default fleet assignment for new Windows MDM enrollments (Fleet Premium), applied before the Autopilot Enrollment Status Page runs so software, scripts, and configuration profiles take effect during out-of-box setup.

— Timing and scope explained, no config key shown `fleet-v4.91.0`

11 **New audit log activities for MFA and setup scripts** NEW — 65

Adds `user_mfa_requested` activity recorded when valid credentials are submitted for an MFA-enabled account and a verification email is sent. Adds `created_setup_experience_script` and `deleted_setup_experience_script` activities to the audit log when a setup experience script is added, replaced, or removed via the API or GitOps.

— Three named activity types, no trigger payload shown `fleet-v4.91.0`

12 **Per-fleet activity automation webhooks** NEW — 60

Adds per-fleet activity automations (Fleet Premium): a webhook configured from the Hosts page that posts to a destination URL whenever an activity tied to one of the fleet's hosts is created.

– UI path given, no payload or endpoint detail `fleet-v4.91.0`

13 **Adobe Creative Cloud plugin inventory** NEW 60

Adds Adobe Creative Cloud plugin detection (CEP and UXP extensions) to software inventory on macOS and Windows, listed as software type `Plugin (Adobe)` with version and host count.

– Scope and labeling named, no config or API surface `fleet-v4.91.0`

THINNER COVERAGE BELOW

14 **Marketing names for Apple hardware models** IMPROVED 55

Adds marketing name display for Apple devices (macOS, iOS, iPadOS) — the `Hardware model` field now shows names like `MacBook Pro (16-inch, 2021)` instead of raw identifiers like `MacBookPro18,1`.

– Before/after example given, no field or API name `fleet-v4.91.0`

15 **Updated macOS CIS benchmark policies** IMPROVED 55

Updates macOS CIS benchmark policies to macOS 14 Sonoma v3.1.0, macOS 15 Sequoia v2.1.0, and macOS 26 Tahoe v1.1.0.

– Exact benchmark versions named, no migration detail `fleet-v4.91.0`

16 **'Patch when closed' patch policy option** NEW 45

Adds a `Patch when closed` option for patch policies that only patches an app when it is not currently running.

– Behaviour stated simply, no config detail `fleet-v4.91.0`

17 **'Added to Fleet' hosts table column** NEW 45

Adds a sortable `Added to Fleet` column to the hosts table showing when each host last enrolled.

– Clear UI location, no further mechanism `fleet-v4.91.0`

18 **Linux platform label option** NEW 40

Adds `linux` as a label platform option, targeting hosts across any Linux distribution.

– Bare one-line addition, no mechanism detail `fleet-v4.91.0`

19 **Managed local admin accounts on Windows** NEW 40

Adds a managed local admin account capability for Windows hosts (requires fleetd 1.60.0 or higher).

– Capability named but no mechanism or config given `fleet-v4.91.0`

20 **Nested Entra group support in IdP vitals** IMPROVED 30

Adds support for nested Entra groups in IdP vitals.

— Bare one-line mention, no mechanism given `fleet-v4.91.0`

WAS THIS USEFUL? 👍 👎

CrowdStrike

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

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

> CrowdStrike introduced coordinated multi-agent SOC investigations built on a new shared-memory layer, a no-code agent builder (AgentWorks) that runs security workflows conversationally inside Anthropic Claude, and new integrations bringing GPT-5.6 Cyber reasoning and OpenAI Codex agent protection into the Falcon platform.

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

01 **Charlotte AI AgentWorks with Claude execution** NEW 69

New Charlotte AI AgentWorks lets security teams describe a desired security outcome in natural language and build, test, and refine custom agents grounded in Falcon telemetry, with no coding required. Security workflows such as triage, enrichment, hunting, and response can then be run conversationally within Anthropic Claude via AgentWorks, with results returned inline and full auditability.

– Describes mechanism and where workflows run, but no exact command. 2026-09-02-8b8eb0b9

THINNER COVERAGE BELOW

02 **Shared context layer across Charlotte AI agents** NEW 56

A new shared context layer, built on Enterprise Graph, gives all Charlotte AI domain agents a persistent, shared memory across every agent, investigation, and tenant — what one agent learns, all agents know — enabling coordinated parallel investigations across the SOC.

– Names Enterprise Graph and mechanism but no UI/API path given. 2026-09-02-8b8eb0b9

03 **Falcon Guardian for OpenAI Codex agents** NEW 42

New Falcon Guardian capability extends CrowdStrike enterprise security to OpenAI Codex agents at the point of execution.

– Names the target (OpenAI Codex) but not how protection works. 2026-09-02-8b8eb0b9

04 **Falcon platform on Anthropic Claude Marketplace** NEW 39

Falcon platform is coming to the Anthropic Claude Marketplace, allowing Anthropic customers to procure it against existing Anthropic commitments.

– States procurement path but no mechanism or timing detail. 2026-09-02-8b8eb0b9

05 **GPT-5.6 Cyber integration into Falcon** NEW 39

GPT-5.6 Cyber, an advanced reasoning model, is being integrated into the Falcon platform, starting with the CrowdStrike Frontier AI Readiness and Resilience (FAIRR) Service.

— Names the model and starting service but no capability detail. 2026-09-02-8b8eb0b9

06 **Charlotte Agentic SOAR workspace** NEW 30

New Charlotte Agentic SOAR provides a single-governed automation workspace for orchestrating multi-agent SOC workflows.

— Names the product but gives no mechanism or scope. 2026-09-02-8b8eb0b9

WAS THIS USEFUL? 👍 👎

Trend Vision One

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

RANK

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

> Trend Vision One added administrator controls and detailed daily digest emails for failed scheduled report notifications.

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

01 **Failed scheduled report notifications** NEW 75

Administrators can enable or disable scheduled report failure notifications via `Administration > Notifications`. By default, a daily failure digest email is sent to the master administrator and the report schedule creator whenever scheduled reports fail, including the report name, report ID, schedule ID, failure time, and failure reason. Notifications also surface clear failure codes and descriptions, widget-level error details for PDF report failures, and configuration-specific guidance such as reminders to verify required frameworks and asset groups.

— Names UI path and email fields but no example shown.

Failed scheduled report notifications (…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

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

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260902.0 introduces a merge-users capability for consolidating duplicate user accounts, gated behind a feature flag and exposed via the API with bypass enforcement.

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

01 **Merge-users capability for account consolidation** NEW

45

Adds a merge-users feature that allows consolidating user accounts, controlled by a feature flag, exposed through an API surface, and subject to bypass enforcement for access control.

— Names feature flag and API surface but no exact flag name or endpoint given. 7.260902.0

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

CrowdStrike Falcon Onum

Feedback ⓘ 1 RELEASE · 2026-09-02

PRESS ↗ RANK

> Falcon Onum adds Certified Data Pipelines that filter noise and run threat detection in-stream before ingested data lands in storage.

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

01 **Certified Data Pipelines for in-stream filtering and detection**

NEW 48

Built on Falcon Onum's streaming pipeline technology, Certified Data Pipelines filter noise at ingestion so only relevant data reaches downstream agents, and run detection inside the pipeline so threats are caught in-stream before ingested data is stored.

— Explains mechanism but no config, flags, or limits given. 2026-09-02-13bacd74

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ 1 RELEASE · 2026-09-02 BLOG ↗ RANK

Tenzir is a data pipeline engine for security teams that collects, transforms, enriches, and routes telemetry.

> Tenzir's v6.15.0 release adds Splunk HEC acknowledgement support for reliable delivery confirmation, new raw Avro decoding primitives, finer-grained parallelism controls, parallelized Google SecOps output, and up to 9x faster event branching.

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

01 **Splunk HEC acknowledgement support** NEW 95

A new `ack = true` option on the `accept_splunk` operator enables the HEC acknowledgement protocol: clients receive an `ackId` per request, queryable via `/services/collector/ack`, with acknowledgements held for 10 minutes by default and delivery guarantees preserved across node crashes/restarts.

Enable HEC acknowledgements on a Splunk ingestion pipeline so clients can confirm delivery and safely retry across node restarts.

```
$ accept_splunk hec_token = secret("splunk-hec-token"), ack =
true
```

— Names exact flag, endpoint, and default timeout with a runnable example

Tenzir Node v6.15.0: Splunk HEC Acknowl…

02 **Expanded parallel operator controls** IMPROVED 75

The `parallel` operator now accepts `fuse` (values: `none`, `parallel`, or `all`) and `limit_partitions` options, matching the `// parallelism` directive and capping keyed operator parallelism for operators like `summarize` and `deduplicate`.

Cap keyed parallelism inside a parallel block to avoid runaway partition growth when summarizing high-cardinality data.

```
$ parallel 8, fuse = "none", limit_partitions = 2 {
  where y > 100
  summarize g, total = sum(x)
}
```

— Names both new options and affected operators with a runnable example

Tenzir Node v6.15.0: Splunk HEC Acknowl…

03 Raw Avro decoding via read_avro and parse_avro NEW

70

The new `read_avro` operator decodes streams of concatenated raw Apache Avro binary datums (not Avro Object Container Files), and the new `parse_avro` function decodes a single raw Avro datum from each `blob` or `string` value; both require the writer schema via the `schema` argument.

Parse raw Avro data from a file and forward to a pipeline — useful when ingesting binary Avro logs from object storage or Kafka.

```
$ read_avro < data.avro | ...
```

— Names both operators and the required schema argument, with example

Tenzir Node v6.15.0: Splunk HEC Acknowl…

THINNER COVERAGE BELOW

04 Parallelized Google SecOps output IMPROVED

35

The `to_google_secops` operator now runs as multiple instances when a pipeline is parallelized, each batching, authenticating, and sending independently for higher-throughput Google SecOps pipelines.

— Names the operator but gives no config example or numbers

Tenzir Node v6.15.0: Splunk HEC Acknowl…

05 Faster event branching with if IMPROVED

35

Pipelines branching events with `if` now achieve up to 9x higher throughput when matching events are contiguous in a batch, by handing views instead of copies for batches with up to four contiguous matching runs.

— Explains mechanism and gives a number but no usable surface

Tenzir Node v6.15.0: Splunk HEC Acknowl…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Yubico

Feedback ⓘ 1 RELEASE · 2026-09-02 BLOG ↗ RANK

Get the YubiKey, the #1 security key, offering strong two factor authentication from industry leader Yubico.

> Yubico's ecosystem saw OpenAI mandate hardware-backed passkeys for its Trusted Access for Cyber program, disabling weaker authentication methods, while the co-branded YubiKey bundle expanded to new markets with two new custom hardware keys.

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

01 Hardware-backed passkey mandate for OpenAI Trusted Access for Cyber

80

BREAKING

OpenAI's Advanced Account Security (AAS) program now requires hardware-backed passkeys for all Trusted Access for Cyber (TAC) members, disabling weaker phishable methods such as SMS OTPs and push notifications upon security key enrollment. Enrolling a YubiKey under the AAS program at `chatgpt.com/advanced-account-security` covers both ChatGPT and Codex accounts with hardware-backed passkey protection.

— Names program, mechanism and exact enrollment URL.

Yubico expands OpenAI partnership to ne…

THINNER COVERAGE BELOW

02 OpenAI + YubiKey bundle expansion and new custom keys NEW

55

The custom 2-pack OpenAI + YubiKey bundle expands to 10 new countries: Australia, Egypt, Japan, Malaysia, Mexico, Saudi Arabia, Singapore, South Korea, Taiwan, and the United Arab Emirates. Two new custom-branded hardware keys are available at preferred pricing for the transition: the YubiKey C NFC – OpenAI (tap-to-authenticate for mobile/tablets) and the YubiKey C Nano – OpenAI (low-profile USB-C for laptops).

— Lists countries and key names but no purchase/action steps.

Yubico expands OpenAI partnership to ne…

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ 2 RELEASES · 2026-09-02 → 2026-09-03 BLOG ↗ RANK

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

> Auth0 shipped Early Access self-service controls letting Organization Admins configure Cross App Access and third-party application access directly via the My Organization API and Embeddable UI Components, alongside a redesigned Dashboard navigation and information architecture.

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

01 Self-service XAA and third-party app access for Organizations

NEW 85

The My Organization API and Embeddable UI Components now expose self-service configuration of third-party application access for an Organization and for individual connections, plus the Cross App Access (XAA) Resource Application role for a connection, so Organization Admins can manage access without routing changes through the Auth0 Dashboard. Connections gain a `third_party_client_access` object (default `{ "default_value": "block", "allowed_values": ["block"] }`) and settings for User Access Authorization and Client Credential Access Authorization, each configurable as Unauthorized, Authorized, or All. XAA configuration carries a domain-verification guardrail — the connection's identity provider must use a domain verified for the associated Organization — and each delegated setting is independently controllable so platform teams choose which settings are exposed to Organization Admins.

Configure Cross App Access (XAA) and user/client credential access authorization for an Organization connection via the Auth0 dashboard.

📍 1. Go to Auth0 Dashboard › Organizations › [your org] › Connections. 2. Select a connection and open its settings. 3. Under Cross App Access (XAA), enable Resource Application for the connection. 4. Set User Access Authorization to Unauthorized, Authorized, or All. 5. Set Client Credential Access Authorization to Unauthorized, Authorized, or All. 6. Save changes.

— Names config object, defaults, and API/UI surfaces; some mechanism on guardrail

My Organization API and Embeddable UI C…

THINNER COVERAGE BELOW

02 Dashboard navigation and IA refresh NEW 50

A new Early Access opt-in redesigns the Auth0 Dashboard sidebar with improved information architecture, consistent breadcrumbs, and a modernized visual experience. Users can toggle the new navigation on or off anytime from the profile menu, or accept it via a `Try now` prompt on login.

Enable the new navigation experience from your profile menu without waiting for a forced rollout.

> 📍 1. Log in to the Auth0 Dashboard. 2. Click `Try now` in the popup that appears on login, OR 3. Open your profile menu in the top-right corner and toggle `New Navigation` on or off at any time.

— Clear opt-in steps given but visual changes described generically

`Dashboard Navigation & IA Refresh i…`

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ **Continuous Control Monitoring (CCM)**

RegScale

Feedback ⓘ `1 RELEASE · 2026-09-03` `BLOG ↗` `RANK`

RegScale is a compliance automation platform that helps organizations manage regulatory requirements, assessments, and evidence collection.

> RegScale CLI 6.44.0 adds a validation check to the JCAM configuration workflow to catch misconfigured attribute mappings.

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

01 **JCAM config validation in CLI** `IMPROVED` 75

The `jcam check_config` command now validates `jcamAgencyDefinedDataItems` mappings against the attributes JCAM actually carries, reporting entries configured for attributes JCAM does not have and attributes JCAM has that are not mapped.

— Names exact command and config field, but no output example shown.

`CLI 6.44.0 (2026-09-03)`

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

LogicGate

Feedback ⓘ 1 RELEASE · 2026-09-02 BLOG ↗ RANK

LogicGate Risk Cloud is the leading AI-powered GRC platform with 30+ applications for enterprise risk, third-party risk, cyber risk and compliance management.

> LogicGate's Summer 2026 release brings its GRC Agents to general availability for four use cases and adds automated control testing, bulk record linking, a lifecycle management environment, four new integrations, and an 81% API performance boost for enterprise customers.

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

01 Automated Control Testing on evidence submission NEW

65

Automated Control Testing now tests linked evidence records automatically as they are submitted, delivering a pass/fail control-level evaluation with detailed rationale and immediate gap flagging instead of waiting for the next manual audit cycle.

— Mechanism and behaviour described, but no configuration steps

Your GRC Agents Are Live: What's Inside…

THINNER COVERAGE BELOW

02 GRC Agents reach general availability NEW

55

GRC Agents are now generally available for four use cases — Third-Party Risk Management (TPRM), Enterprise Risk Management (ERM), AI Governance, and Business Continuity Management (BCM) — handling high-volume, repetitive workflow administration so risk teams can focus on proactive risk management.

— Names four use cases but no setup steps given Your GRC Agents Are Live: What's Inside…

03 SmartSheet, IronClad, Box and Microsoft Planner integrations

NEW

55

SmartSheet, IronClad, Box, and Microsoft Planner join the LogicGate integration ecosystem for connected document management, task tracking, and contract execution.

— Names all four integrations but no config detail

Your GRC Agents Are Live: What's Inside…

04 Provision 29 Compliance application co-developed with PwC NEW

55

Provision 29 Compliance application — co-developed with PwC — provides UK enterprise risk teams a single, connected workflow for audit-ready principal risk and material

control management evidence.

— Names app and partner, but no usage steps Your GRC Agents Are Live: What's Inside…

05 **GRC Lifecycle Management with dev sync** NEW 50

GRC Lifecycle Management lets teams sync a live production instance to a dedicated development environment to safely test and build new workflows without risking live processes.

— Describes mechanism but not activation steps Your GRC Agents Are Live: What's Inside…

06 **MCP Server Access for natural-language queries** NEW 50

Upcoming MCP Server Access will allow users to query Risk Cloud data via natural language prompts in the LLM of their choice, with permission-based access and no custom integration required.

— Upcoming feature, not yet available to act on Your GRC Agents Are Live: What's Inside…

07 **Upcoming Insights Agent and additional GRC Agents for fall** NEW 50

Upcoming Insights Agent will power next-generation GRC reporting; additional GRC Agents for Cyber Risk Management, Policy & Procedure Management, Incident Management, and Issues Management are confirmed for fall.

— Roadmap item naming four future agents, nothing usable yet

Your GRC Agents Are Live: What's Inside…

08 **Bulk Record Linking at scale** NEW 45

Bulk Record Linking applies the same semantic intelligence behind Record Linking Recommendations at scale, generating recommendations for hundreds of records simultaneously rather than one at a time.

— Scale and mechanism named, no usage detail Your GRC Agents Are Live: What's Inside…

09 **API performance upgrades cut record load times** IMPROVED 40

API performance upgrades deliver an 81% reduction in record load times for enterprise-sized customers.

— Concrete number given but no mechanism detail Your GRC Agents Are Live: What's Inside…

10 **Upgraded Search across applications, reports and dashboards** IMPROVED 35

Upgraded Search now extends beyond individual records to applications, reports, and dashboards.

— Names scope but no mechanism or path Your GRC Agents Are Live: What's Inside…

11 **Visual Reports branding and branded email sending** IMPROVED

35

Visual Reports gain full color-palette customization to match corporate branding, and automated emails can now be sent from an organization's own domain.

— Two thin additions named, no navigation detail Your GRC Agents Are Live: What's Inside…

12 **Spark AI foundational skills built into platform** IMPROVED

35

Foundational AI skills previously branded as Spark AI are now built directly into the Risk Cloud platform with no separate product layer to manage.

— Describes consolidation but not new capability specifics

Your GRC Agents Are Live: What's Inside…

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

CrowdStrike Falcon Secure AI Services

Feedback ⓘ

1 RELEASE · 2026-09-02 PRESS ↗ RANK

CrowdStrike Falcon Secure AI Services provides AI-powered threat detection, incident response automation, and security intelligence for enterprise endpoints and cloud environments.

> CrowdStrike introduced the FAIRR Service, integrating OpenAI's GPT-5.6 Cyber model into its cyber harness to assess frontier AI risk and attack paths.

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

01 **FAIRR Service with GPT-5.6 Cyber** NEW 45

New FAIRR Service applies GPT-5.6 Cyber within CrowdStrike's cyber harness to assess frontier AI risk, analyze attack paths, and inform remediation priorities. It is available for approved, authorized defensive use cases with expert oversight, per the CrowdStrike/OpenAI partnership announcement.

— Names the service and model but access/mechanism vague, no how-to. 2026-09-02-2a585013

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 3, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue