THE CYBER TOOLCHAIN

NO. 065

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 5, 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 71 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 23 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. Aurascape stops secrets and destructive commands inside a live Claude Code or Cursor session,

vim ? help

before any artifact exists to scan. GitGuardian now runs its full detection engine against vaulted secrets, classifying them by what they actually are. And Arnica names every human and service account with code access across SCMs.

PROTECT

Catch a bad AI coding action while the session is running, not after the commit

Aurascape inspects live sessions in Claude Code, GitHub Copilot and Cursor and blocks secrets, destructive commands and prompt injection at the moment they happen — conventional scanning has nothing to look at until a file or commit exists. Aikido's Deep Review works the other end, reasoning across the whole codebase on a PR and remembering earlier review decisions so the same finding isn't re-litigated every cycle.

Aurascape · Aikido Security

IDENTIFY

Know what's really in a vault entry, and who can reach your code

GitGuardian scans ggscout-collected vaulted secrets with its full Secrets Detection Engine and classifies by actual value rather than the name or path someone typed years ago, so a key mislabeled as a config string stops hiding. Arnica's cross-SCM identity inventory answers the question most orgs cannot answer on demand: every person and service account that can push code, in one place.

GitGuardian SaaS · Arnica

GOVERN

Route and cap LLM traffic through one control point

The gateway puts guardrails, usage limits and model access controls in front of multi-workspace LLM traffic, plus SAML `No access` roles and default-deny project visibility. That turns AI usage policy into something enforced at the routing layer rather than documented in a wiki and hoped for.

Openlayer

PROTECT

Suppress a finding without it disappearing quietly

Ignored findings can now be gated behind an approval workflow in public beta, so a developer muting a result creates a reviewable decision instead of a silent hole. Read-only tokens and the `is_partial_scan` API

> filter also let you wire dashboards and automation without handing out write access or mistaking a partial scan for full coverage.
>
> Semgrep

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ 1 RELEASE · 2026-09-05 NOTES ↗ RANK

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

> Strix v1.6.2 adds a new web search provider for agent reconnaissance and sharpens error telemetry with more granular classification.

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

01 **Exa web search provider for agents** NEW 38

Strix now supports Exa as a web search provider alongside Perplexity, giving agent-driven reconnaissance an additional search backend to draw on.

— Names the provider but no config/flag to select it v1.6.2

02 **Phase and exception-class error telemetry** IMPROVED 35

Error beacons are now classified by phase and exception class, giving richer operational diagnostics for troubleshooting Strix runs.

— Describes the change but no surface a reader can act on v1.6.2

WAS THIS USEFUL? 👍 👎

T3MP3ST

Feedback ⓘ changes since 2026-09-04 CODE ↗

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

T3MP3ST adds a WPScan JSON parser, known-debt adapter templates, a format-string CTF lab, and a synthetic memory-forensics fixture.

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-09-04 https://github.com/elder-plinius/T3MP3ST.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-04
```

└──▷ **TRY IT**

Regenerate the memory-forensics fixture byte-for-byte to verify provenance before running the CTF challenge.

```
$ python3 ctf/challenges/artifacts/memory-forensics/generate.py
```

Adds `parseWpscan` to `src/arsenal/parsers.ts`, ingesting `wpscan --format json` output and mapping WordPress core, plugin, and theme vulnerabilities to structured `ToolFinding` records with severity levels (high for core `insecure` labels, medium for confirmed component vulns, info for interesting findings).

Adds `ARG_TEMPLATES` for five previously `KNOWN_DEBT` adapters — `feroxbuster` (`-u` / `-w`, `--json -q`), `osv-scanner` (`--format json --recursive <path>`), `hashcat` (receipt-gated with `--potfile-disable`), `apktool` (`d -f <apk> -o apk-out`), and `yara` (explicit ruleset required) — emptying the invocation-honesty guard's `KNOWN_DEBT` set.

Adds a reproducible native format-string CTF lab (`ctf/docker/pwn/format-string/vuln.c`) exposing an intentional `dprintf` format-string write primitive on port `9002`, with digest-pinned images, loopback-only publish, and a deterministic solver.

Adds a deterministic synthetic memory-forensics fixture (`ctf/challenges/artifacts/memory-forensics/memdump.raw`) generated from `generate.py`, using `FORMAT_VERSION = 'T3MP3ST-SYNTH-MEM-v1'`, a fixed 32 768-byte size, and SHA-256-seeded deterministic noise — replacing an opaque host-mounted dump with a fully reproducible offline artifact.

WAS THIS USEFUL?  

Hacktron

Feedback

1 RELEASE · 2026-08-19

NOTES ↗

RANK

Hacktron is an AI-powered security code-review platform. Our AI agents review your repositories for exploitable vulnerabilities and surface them as pull-request comments and dashboard findings.

> Hacktron split billing into separate Automation and Whitebox Scan credit accounts, expanded on-premises deployments with source archive uploads and air-gapped Git support, and added Microsoft account linking plus glob-based author filters.

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

01 On-premises Whitebox Scan source options NEW

75

On-premises deployments can now select enabled GitHub Enterprise Server repositories or upload source archives in `.zip`, `.7z`, `.tar.gz`, and `.tgz` formats when starting a Whitebox Scan.

Run a Whitebox Scan on an air-gapped codebase by uploading a source archive directly from the on-premises deployment UI.

> 📍 When starting a Whitebox Scan in the on-premises deployment, choose Upload source archive and select your prepared `.tar.gz` (or `.zip`, `.7z`, `.tgz`) file containing the repository source.

— Lists exact archive formats plus a runnable upload example

Separate credit accounts and more on-pr…

02 Microsoft account linking for Azure DevOps attribution NEW

70

Users can link a Microsoft account so Azure DevOps pull request activity is attributed to the correct Hacktron member and counts toward their Code Review seat.

Link a Microsoft account so Azure DevOps pull request activity is attributed to the correct Hacktron member and counts toward their Code Review seat.

> 📍 In Hacktron, go to Personal settings › Connected accounts and select Connect Microsoft account, then sign in with the Azure DevOps identity whose UPN matches your pull request author email.

— Includes exact navigation path via Personal settings > Connected accounts

Separate credit accounts and more on-pr…

03 Glob patterns in author filters IMPROVED

65

Repository and organization author filters now accept glob patterns (`*` , `**` , `?` , `{a,b}`) in addition to exact names, with bounded limits on pattern length, wildcard count, and brace alternatives.

— Names exact glob syntax and limits but no config file location

Separate credit accounts and more on-pr…

04 **Separate Automation and Whitebox credit accounts** NEW

63

Billing now shows separate accounts for Automation credits and Whitebox Scan Credits, each with its own balance, current-period usage, purchase flow, and history. Auto top-up is available for Automation credits only.

— Names the two accounts and auto top-up scope but no UI path

Separate credit accounts and more on-pr…

THINNER COVERAGE BELOW

05 **Air-gapped appliance connectivity controls** NEW 50

Appliance operators can allow specific private Git hosts, add trusted CA certificates, and collect masked support bundles without an outbound internet connection, enabling air-gapped support.

— Describes capabilities but no exact config path or command

Separate credit accounts and more on-pr…

06 **Automation deletion and incomplete-config safeguards** IMPROVED

35

Deleting an automation now shows a confirmation dialog with its name, and attempting to enable an incomplete automation identifies the missing configuration.

— Describes UX behaviour only, no named fields or commands

Separate credit accounts and more on-pr…

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ 1 RELEASE · 2026-09-05 CHANGELOG ↗ RANK

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

> Semgrep rolled out a public beta for gating ignored findings behind an approval workflow, added read-only tokens, multi-webhook remediation policies, and a new `is_partial_scan` filter on the scans API, alongside smaller platform UI changes.

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

01 **Developer exception requests for ignored findings** NEW 50

A public beta lets admins gate ignored findings behind an approval workflow by adding Require exception requests to any policy, so developers must request an exception rather than silently ignoring a finding.

– Names the policy toggle but no approval flow mechanism detailed

changelog-20260905-da089ddc

02 **Separate Usage and Billing tabs in platform settings** IMPROVED 50

AppSec Platform Settings splits Usage and Billing into separate tabs; credit usage is now visible to all users, and self-serve product purchasing is available.

– Names exact tabs and access change but no further detail changelog-20260905-da089ddc

03 **Multiple webhooks per remediation policy** IMPROVED 35

Remediation policy automations now support adding more than one webhook action to a single policy, rather than being limited to one.

– Clear before/after but no config surface or steps given changelog-20260905-da089ddc

04 **Read-only API tokens** NEW 10

Semgrep adds read-only tokens, per the release summary.

– Only named in summary, no mechanism or usage described changelog-20260905-da089ddc

WAS THIS USEFUL? 👍 👎

Snyk

Feedback ⓘ 1 RELEASE · 2026-09-05 DOCS ↗ BLOG ↗ RANK

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

> Snyk shipped Snyk Assist, an AI-powered assistant embedded in the new Snyk UI that answers product and vulnerability questions, retrieves account data, and can file support cases without leaving the product.

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

01 **Snyk Assist AI assistant in the UI** NEW 76

Snyk Assist is an AI-powered assistant accessible from the top right of every page in the new Snyk UI, answering product, setup, and troubleshooting questions grounded in Snyk documentation, release notes, and support knowledge articles; it is available on Team, Ignite, and Enterprise plans. It can retrieve read-only account information scoped to the current organization and group using existing permissions — issue summaries by severity, issue triage and single-issue detail, projects and import targets, collections and container images, scan settings, integrations and their connectivity status, members and their roles, and a project's current ignore rules. It can also look up known vulnerabilities for a package or specific package version against the Snyk Vulnerability Database, and can suggest and create a support case from within the conversation (returning the case ID and link) or capture unsupported capabilities as feature requests.

When onboarding a new team and unsure which GitHub permissions Snyk requires, ask Assist in plain language to get a documentation-grounded answer without leaving the product.

> 📍 In the Snyk UI (new navigation), click the Snyk Assist icon in the top right corner, then type: 'What permissions does Snyk need for my GitHub org?'

When you need to audit which integrations are connected and their status for your current org, query Snyk Assist instead of navigating through settings screens.

> 📍 In the Snyk Assist panel, type: 'Show me the integrations and their connectivity status for my current organization.'

When triaging a dependency and needing quick vulnerability data, ask Snyk Assist to look up known CVEs for a specific package version against the Snyk Vulnerability Database.

> 📍 In the Snyk Assist panel, type: 'What known vulnerabilities exist for lodash version 4.17.20?'

— Full capability list and mechanism given; only UI path, not a command, for access.

2026-09-05-5fea215a

WAS THIS USEFUL? 👍 👎

Aikido Security

Feedback ⓘ 1 RELEASE · SEEN 2026-09-05 NOTES ↗

RANK

Aikido Security is a developer-focused platform that identifies and helps remediate security vulnerabilities in code and infrastructure.

> Aikido launches Deep Review, an agentic pull request reviewer that reasons across full codebases to catch deep security issues and remembers past review decisions, priced per PR from a spend-capped wallet.

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

01 **Deep Review agentic PR reviewer** NEW 85

Deep Review is an agentic pull request reviewer that reasons across your full codebase — including middleware, global config, and linked repositories — to catch deeply nested security issues, broken access control, business logic flaws, and cross-tenant data leaks, posting findings with suggested fixes directly on PRs. It saves accepted-risk and false-positive decisions to memory so future reviews on the same repository learn from prior feedback, and runs after existing SAST and dependency checks to build on their findings rather than repeat them. It is enabled and configured per repository, with linkable related repos and a configurable review level, and is priced at 1 credit ($1) per PR from the Aikido Wallet with unlimited re-reviews on the same PR and a configurable monthly spend cap.

Enable Deep Review on a repository and link a related repo so it reasons across both codebases when reviewing PRs.

> 📍 In the Aikido console, go to your repository settings, enable Deep Review, link related repositories under Linked repos, and set the review level that fits your team.

Keep monthly spend predictable by setting a cap on how many credits Deep Review can consume across your repos.

> 📍 In the Aikido console, open Aikido Wallet settings and set a monthly credit cap to limit Deep Review spend across all enabled repositories.

— Explains mechanism, scope, pricing and config, with UI steps but no CLI/API.

🚀Deep Review: the agentic PR review th…

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ 3 RELEASES · 2026-09-04 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 shipped GPT-6 Astra as a generally available Copilot model across every major IDE and surface, alongside new Claude Fable 5.1 and Gemini 3.8 Flash model options, a public preview of Agent Merge in VS Code, and a new privacy-safe star history REST API endpoint.

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

01 **GPT-6 Astra generally available in Copilot** NEW 90

GPT-6 Astra, OpenAI's agentic long-horizon coding model, is now selectable in the Copilot model picker across VS Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, the Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse. Copilot Enterprise and Business admins can enable or restrict access via the default model enablement control under the model policy in Copilot settings.



— Names every supported surface and the exact admin control

GPT-6 Astra is generally available in G…

02 **Agent Merge public preview in VS Code** NEW 60

Agent Merge is now in public preview in VS Code 1.136, resolving review feedback, failed checks, and merge conflicts to get pull requests merge-ready.

— Names the VS Code version and what it resolves GitHub Copilot weekly releases — August…

THINNER COVERAGE BELOW

03 **New Copilot models: Claude Fable 5.1 and Gemini 3.8 Flash**

NEW

50

Copilot adds Claude Fable 5.1 for Pro+, Max, Business, and Enterprise users, and Gemini 3.8 Flash rolling out to Pro, Pro+, Max, Business, and Enterprise users.

— Names models and tiers but no usage mechanism GitHub Copilot weekly releases — August…

04 **VS Code chat and agent workspace enhancements** NEW

45

VS Code's Copilot chat and Agents experience gains three additions: experimental multi-root workspaces that bring Copilot and Claude agent sessions to every folder in a workspace, experimental chat backgrounds that let you personalize the Agents window with built-in patterns or custom images, and chat sessions that now organize related chats into a hierarchy and surface which ones need attention.

— Lists three named UI additions without deeper mechanism

GitHub Copilot weekly releases — August…

05 **Privacy-safe star history API endpoint** NEW

45

A new REST API endpoint returns a repository's historical star counts with timestamps without exposing individual stargazer identities.

— Describes behavior but omits exact method and path

New API endpoint provides privacy-safe …

06 **Copilot harness generally available for JetBrains** IMPROVED

25

The GitHub Copilot harness is now generally available in Copilot for JetBrains, enabling faster feature delivery and better code quality.

— Bare GA announcement with vague benefit claim GitHub Copilot weekly releases — August…

WAS THIS USEFUL?  

◆ Web Application Security

caido

Feedback ⓘ 1 RELEASE · 2026-09-04 NOTES ↗ RANK

Caido releases, wiki and roadmap

> Caido v0.58.3 improves the request editing workflow with inline autocomplete and tighter integration between Replay and pipeline sessions.

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

01 **Add existing Replay requests to pipeline sessions** IMPROVED

40

Users can now add existing Replay requests directly to a pipeline session, rather than only creating new ones.

— Describes the workflow change but lacks UI path or command detail v0.58.3

02 **Inline autocomplete in HTTP request editor** IMPROVED

30

The HTTP request editor now shows inline autocomplete hints while typing requests.

— Names the feature but no mechanism or scope detail v0.58.3

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ 1 RELEASE · 2026-09-04 BLOG ↗

RANK

GitGuardian SaaS scans git repositories and code to detect and prevent secrets, credentials, and sensitive data from being exposed.

> GitGuardian now runs its full Secrets Detection Engine against vaulted secrets collected by ggscout, classifying them by actual value rather than name or path, and adds finer-grained per-team incident permission controls in Teams Management.

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

01 Secrets Detection Engine classification for vaulted secrets

NEW 85

Requires `ggscout` version `0.32.0` or later. Runs the Secrets Detection Engine — the same 550+ detectors used for code scanning, including multi-match detectors — against every value `ggscout` collects from vaults, CI systems, and Kubernetes clusters, classifying each secret by its actual value as a real credential, noise, or `Unknown` rather than by its name or path. Adds a `Detector` column and filter controls to the NHI Governance → Identities inventory, filterable by `detector`, `detector type`, `secret family`, `secret category`, and `secret provider` — the same filter set used for code-scanning incidents. Groups sibling fields that form a single credential (e.g. PostgreSQL username, password, host, and port) into one inventory entry instead of four separate items, and retains unmatched secrets labeled `Unknown` so they stay visible and filterable rather than being silently dropped.

Filter your NHI inventory down to real credentials by detector type and secret provider, excluding noise and unknowns.

> 📍 In the NHI inventory, open the Detector column filter and select the desired `detector type`, `secret family`, `secret category`, or `secret provider` to isolate real credentials from placeholders and feature flags.

Review classified vault secrets filtered to a specific provider in the NHI inventory after ggscout has run.

> 📍 In the GitGuardian console, go to NHI Governance → Identities, click the Detector column header to sort, then open the filter panel and select a value under `Secret provider` to scope the inventory to credentials from that provider.

— Names detectors, columns, filters, grouping mechanism and requirement, but only UI path given

Vaulted Secrets Now Classified by Value…

THINNER COVERAGE BELOW

02 **Per-team incident permissions in Teams Management** NEW

36

When promoting a user to Member, admins can now set per-team incident permissions, apply one permission in bulk across teams, and edit team rights after promotion.

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

Vaulted Secrets Now Classified by Value…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL?  

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

quark-engine

Feedback ⓘ 1 RELEASE · 2026-09-04 NOTES ↗ RANK

Quark-Engine is a mobile security analysis framework that detects malicious behaviors in Android applications through rule-based pattern matching.

> Quark-Engine made DexTrace the default core analysis library for APK scanning, requiring a flag to restore the previous Androguard-based behaviour.

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

01 **DexTrace as default core library** BREAKING 75

The `--core-library` flag now accepts `[androguard|rizin|radare2|shuriken|dextrace]`, and `dextrace` is the new default for APK analysis, replacing Androguard. Workflows relying on the previous Androguard-based default behaviour must explicitly pass `--core-library androguard` to preserve it, e.g. `quark --core-library androguard -a <sample.apk>`.

Analyse an APK with the legacy Androguard library when DexTrace produces different results for your sample.

```
$ quark --core-library androguard -a <sample.apk>
```

— Names exact flag, values, and fallback command; lacks internal mechanism detail.

v26.9.1

└──▷ BREAKING ON UPGRADE

! The default analysis library is now DexTrace instead of Androguard; workflows that relied on Androguard behaviour by default must add `--core-library androguard` to preserve the previous behaviour.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

2 RELEASES · 2026-08-24 → 2026-08-27 BLOG ↗ RANK

Finite State Platform is a software supply chain security platform that analyzes and monitors firmware and software for vulnerabilities and compliance risks.

> Finite State Platform expanded firmware unpacking to ESP32 and Airoha Bluetooth-audio images, upgraded binary scan SBOM exports to CycloneDX 1.6 with author metadata, and added a License Policy setting for components with no license.

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

01 **License Policy for unlicensed components** NEW 70

Adds a License Policy setting for Components with no license on the Bill of Materials, letting you assign a severity and a label; the label then appears on those Components in the Bill of Materials.

Flag unlicensed components during a compliance review by configuring a severity and label for them in License Policy.

> 📍 In the Finite State Platform, go to License Policy, add a rule for Components with no license, and set a severity and label — the label will then appear on those Components in the Bill of Materials.

— Includes step-by-step config path and effect on BOM

August 27, 2026 - Empty-license policy …

02 **ESP32 and Airoha Bluetooth-audio firmware unpacking** NEW 65

Unpacks Espressif ESP32 firmware images, including full flash dumps, so those devices are analyzed as firmware rather than as an unidentified blob. Also unpacks Airoha Bluetooth-audio firmware, identifying the vendor SDK and DSP components and producing a complete SBOM for those images.

— Names devices and components but no UI/command path given

August 24, 2026 - ESP32 and Bluetooth-a…

THINNER COVERAGE BELOW

03 **Reachability results applied before binary assessment**

IMPROVED 50

Saves reachability results with the scan and applies them before binary assessment so later analysis uses a consistent verdict.

— Explains sequencing effect but no concrete surface named

August 24, 2026 - ESP32 and Bluetooth-a…

04 **CycloneDX 1.6 SBOM export with author metadata** IMPROVED

45

Exports binary scan SBOMs as `CycloneDX 1.6` with author metadata included.

— Names format version but no mechanism or steps August 24, 2026 - ESP32 and Bluetooth-a…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sophos Email Security

Feedback ⓘ 1 RELEASE · 2026-09-05

CHANGELOG ↗ RANK

Sophos Email Security is a cloud-based email protection solution that defends against phishing, malware, ransomware, and spam threats.

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

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

01 Automatic quarantine release via Intelix scanning NEW

63

Emails held in quarantine are automatically released when an Intelix scan determines the email is not malicious or suspicious, removing the need for manual review. Enabled via Sophos Central under Email Security > Policies > Scan Attachments with Intelix by turning on Automatically release from quarantine.

Enable automatic quarantine release so emails cleared by Intelix are freed without manual review, reducing analyst workload.

📍 In Sophos Central, go to Email Security › Policies › Scan Attachments with Intelix, enable Automatically release from quarantine for emails Intelix determines are not malicious or suspicious.

— Names exact UI path and toggle, but no scope/limits detail changelog-20260905-903c0107

THINNER COVERAGE BELOW

02 Configurable notification recipients for Data Control NEW

26

Adds the ability to configure notification recipients for Data Control rule actions.

— Bare description with no mechanism or config path given changelog-20260905-903c0107

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Mobile

Feedback ⓘ 1 RELEASE · 2026-09-05 CHANGELOG ↗

RANK

Sophos Mobile is an enterprise mobile device management and security platform protecting iOS and Android devices with threat defense and compliance controls.

> Sophos Mobile added support for Apple declarative device management (DDM) policies, letting managed iOS devices self-enforce compliance and report status changes in real time.

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

01 Apple declarative device management (DDM) policies NEW

50

Sophos Mobile now supports Apple declarative device management (DDM) policies, enabling desired-state configuration where devices self-enforce compliance and report changes immediately rather than waiting for a periodic sync.

– Explains mechanism but no UI path or config specifics given.

changelog-20260905-ab501c3c

WAS THIS USEFUL? 👍 👎

Sophos Anti-Virus for macOS

Feedback ⓘ

1 RELEASE · 2026-09-05 | CHANGELOG ↗ | RANK

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

> Sophos Anti-Virus for macOS 2026.2.0 adds support for the upcoming macOS 27 Golden Gate release and updates its threat detection engine.

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

01 **macOS 27 Golden Gate support** NEW 30

Sophos Anti-Virus for macOS 2026.2.0 adds support for macOS 27 Golden Gate.

— Names the OS version but no mechanism or scope detail changelog-20260905-066bb8fa

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

The threat detection engine is updated to version 3.99.0.

— Names version number but no detail on what changed changelog-20260905-066bb8fa

WAS THIS USEFUL? 👍 👎

Sophos Central

Feedback ⓘ 1 RELEASE · 2026-09-05 CHANGELOG ↗

RANK

Sophos Central is a cloud-based endpoint protection and management platform that delivers security, threat detection, and device management across organizations.

> Sophos Central adds on-premises Active Directory support to ITDR and rolls out guided Zero Trust workflows for securing SaaS apps, alongside redesigned Identity Overview, Posture Check, and threat protection policy pages.

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

01 **Posture Check page redesign with filtering and search**

IMPROVED 70

The Posture Check page is redesigned with filtering, sorting, and searching by category, provider type, and status, plus tag chips, a live match counter, and a NEW badge for recently published checks.

— Names specific UI controls added, but only a UI path to use them.

changelog-20260905-416580f4

02 **Guided ZTNA workflows for SaaS apps via Protected Browser**

NEW 70

New guided Zero Trust Network Access workflows help secure SaaS apps via Protected Browser, including data boundaries, device posture checks, and in-browser RDP and SSH access, available under the Workspace Protection license.

— Names license and access methods but no exact setup steps. changelog-20260905-416580f4

THINNER COVERAGE BELOW

03 **On-premises Active Directory support in ITDR** NEW 55

Sophos ITDR now supports on-premises Active Directory (AD) via a new sensor deployed to an AD server, extending identity monitoring and posture assessments beyond cloud identity sources.

— Names the sensor mechanism but not deployment steps or limits.

changelog-20260905-416580f4

04 **Redesigned Identity Overview dashboard** IMPROVED 50

The Identity Overview dashboard is redesigned with new widgets covering risk trends, MFA coverage, dormant accounts, credential leaks, and VIP users.

— Lists widgets added but no navigation or config detail. changelog-20260905-416580f4

05 **Updated threat protection policy page** IMPROVED 45

The threat protection policy page is rolled out with a clearer layout, filters, and the option to add explanatory notes to settings and exclusions.

— Describes layout changes generically without specific settings named.

changelog-20260905-416580f4

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ 1 RELEASE · 2026-09-04 NOTES ↗ RANK

Manage devices and everything they touch. See how security, IT, and compliance get easier when unified around devices with Iru.

> Iru shipped directory sync scheduling controls in Iru Identity, new Android work profile account management options, and hardware health visibility for Apple mobile devices.

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

01 **Customizable directory sync schedules and on-demand sync** NEW

35

Iru Identity now supports customizable directory sync schedules and the ability to trigger sync on demand, rather than relying solely on a fixed automatic interval.



– Describes capability but no exact schedule options or UI path

What's New in Iru | September 4th

02 **Managed Google account controls for Android** NEW 25

Android work profiles gain managed Google account support and credential manager controls.

– Names the feature but no mechanism or configuration detail

What's New in Iru | September 4th

03 **Hardware health data for iOS devices** NEW 25

Iru adds hardware component health data visibility for iPhone and iPad devices.

– States the capability without specifics on which components or metrics

What's New in Iru | September 4th

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

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

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

> Google SecOps added self-service access to the Bindplane Enterprise license key directly from the platform console, removing the need to file a support ticket for Enterprise Plus and GUS customers.

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

01 **Self-serve Bindplane Enterprise license key** NEW 80

Enterprise Plus and GUS customers can now download their Bindplane Enterprise (Google Edition) license key directly from the console under `SIEM Settings > Collection Agents`, instead of requesting it through support. This is a Preview feature currently available in US and EU regions only.





[redacted]

Retrieve your Bindplane Enterprise (Google Edition) license key without opening a support ticket — useful when spinning up new collection agents.

> 📍 In the Google SecOps console, go to SIEM Settings › Collection Agents, then locate and download your Bindplane Enterprise (Google Edition) license key.

— Names exact UI path and scope, but limited mechanism detail beyond access change.

September 03, 2026 (2026-09-03)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

SailPoint

Feedback ⓘ 1 RELEASE · 2026-09-04 BLOG ↗ RANK

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

> SailPoint's Delimited File Connector gained certificate-based authentication for SFTP and FTPS sources, and Identity Security Cloud added 11 new connectors for compliance coverage.

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

01 SSH key authentication for SFTP in Delimited File Connector

IMPROVED 93

Adds a Public Key option to the Authentication Type dropdown when File Transport is set to SFTP, enabling SSH private key authentication via the SSH Private Key field (accepting PEM, OpenSSH, or OpenSSH v1 format) with an optional SSH Private Key Passphrase. Keys and passphrases are stored as encrypted attributes in Identity Security Cloud source configuration rather than uploaded to the Virtual Appliance filesystem, and Test Connection now validates the SSH key auth (invalid keys, wrong passphrases, network errors) before aggregation is scheduled.

Configure a Delimited File source to authenticate to a hardened SFTP server using an SSH private key instead of a password.

> 📍 In the Identity Security Cloud admin console, open the Delimited File source › Connection Settings. Set File Transport to SFTP, then set Authentication Type to Public Key. Enter the Host, Port, and Username (SSH account). Paste your PEM or OpenSSH private key into SSH Private Key. If the key is passphrase-protected, enter it in SSH Private Key Passphrase. Click Test Connection to confirm before saving.

— Exact fields and validation behavior, with runnable steps

Enhancement: Delimited File Connector -…

02 TLS client certificate authentication for FTPS in Delimited File Connector

IMPROVED 93

Adds a Client Certificate option to the Authentication Type dropdown when File Transport is set to FTPS, enabling TLS mutual authentication via a Base64-encoded PKCS12 (.p12) certificate pasted into FTPS Client Certificate, with an optional FTPS Client Certificate Passphrase. The certificate and passphrase are stored as encrypted

attributes in source configuration, and Test Connection validates the FTPS TLS handshake and FTP login before aggregation is scheduled.

Configure a Delimited File source to authenticate to an FTPS server using a TLS client certificate (hybrid model: certificate for the TLS layer, credentials for the FTP layer).

> 📍 In the Identity Security Cloud admin console, open the Delimited File source › Connection Settings. Set File Transport to FTPS, then set Authentication Type to Client Certificate. Enter Host, Port, Username, and Password (FTP account). Paste the Base64-encoded PKCS12 (.p12) content into FTPS Client Certificate. If the keystore is protected, enter the passphrase in FTPS Client Certificate Passphrase. Click Test Connection to validate the TLS handshake and FTP login before saving.

— Exact fields and validation behavior, with runnable steps

Enhancement: Delimited File Connector -…

THINNER COVERAGE BELOW

03 **11 new connectors for compliance integration** NEW 35

Adds the AlertOps, Bill4Time, Canny, Documo, Honeybadger, Huntress, Kobiton, Labelbox, PactFlow, Recruitee, and Scalefusion connectors for compliance integration.

— Names all connectors but no setup or mechanism detail product docs

WAS THIS USEFUL? 👍 👎

1Password

Feedback

1 RELEASE · 2026-09-03

BLOG ↗

RANK

All the latest releases from the 1Password team

> 1Password SaaS Manager adds per-API-key cost attribution for OpenAI, Google Gemini provisioning, configurable duplicate access policies, and fixes New Relic user import pagination, while requiring a reconnect for Greenhouse users on the old API version.

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

01 **Greenhouse integration migrated to API v3** BREAKING

50

The Greenhouse integration has migrated to API v3; users who originally connected via API v2 must reconnect Greenhouse using OAuth.

– Clear migration note with required reconnect action

1Password SaaS Manager #1515 (2026-09-0…

02 **Per-API-key cost attribution for OpenAI** NEW

40

1Password SaaS Manager now adds per-API-key cost attribution for OpenAI Platform when the provider supplies API-key cost data.

– Names the surface but no detail on how attribution is surfaced

1Password SaaS Manager #1515 (2026-09-0…

03 **Duplicate access policies with distinct request durations**

NEW

40

Administrators can now create duplicate access policies that have different request durations.

– Names the admin capability but no config surface or limits

1Password SaaS Manager #1515 (2026-09-0…

04 **New Relic user import pagination fix** IMPROVED

40

New Relic user imports now follow pagination, including users beyond the first 500.

– Gives concrete limit fixed but no further mechanism

1Password SaaS Manager #1515 (2026-09-0…

05 **Google Gemini user provisioning** NEW

30

Adds support for provisioning users in Google Gemini.

– Bare statement of new provider support, no mechanism given

1Password SaaS Manager #1515 (2026-09-0…

└──▷ **BREAKING ON UPGRADE**

! Greenhouse integration migrated to API v3; users who originally connected via API v2 must reconnect Greenhouse using OAuth.

WAS THIS USEFUL? 👍 👎

◆ **Data Security & DLP**

Zama fhEVM

Feedback ⓘ 1 RELEASE · 2026-09-04 NOTES ↗ RANK

FHEVM, a full-stack framework for integrating Fully Homomorphic Encryption (FHE) with blockchain applications

> fhEVM v0.13.4 backports Polygon CLI support to the stable 0.13.x branch and updates the relayer to return the full set of user-decrypt shares per request.

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

01 **Relayer returns all user-decrypt shares** IMPROVED 45

The relayer now returns all user-decrypt shares per request, rather than a subset as before.

— States before/after behavior but no config or command v0.13.4

02 **Polygon CLI backported to 0.13.x line** IMPROVED 40

The Polygon CLI has been backported to the 0.13.x release line, enabling Polygon-targeted workflows on the stable branch.

— Names the surface but no usage detail or commands v0.13.4

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ 1 RELEASE · 2026-09-04

NOTES ↗ RANK

The Universal Governance, Risk, Compliance (GRC) Operating System with Integrated Security for Agentic AI, Non-Human Identities, and Swarm Governance. AI SAFE² + AI Sovereignty Maturity Model (AISM), NEXUS-A2A Protocol, FORGE-Act, Marshal Plan for AI [Dual License: MIT + CC-BY-SA]

> AI SAFE² v3.1 introduces a unified `safe2` CLI, structured AISM Decision Cards, new evidence-source integrations with NEXUS and NVIDIA SkillSpector, and a 27-requirement Unbiased AI Standard regulatory profile.

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

01 **Unified `safe2` CLI for scanning and verification** NEW

80

Adds the `safe2` CLI with subcommands including `safe2 scan project` and `safe2 example verify` for scanning AI implementations, evaluating evidence, and validating packaged examples, intended for use in CI/CD pipelines and automated agent workflows.

— Names exact runnable subcommands, no output examples shown

2026-09-04_Agent_CLI_AISM_Decision_Supp…

02 **Unbiased AI Standard (UAS) regulatory profile** NEW

65

Adds the Unbiased AI Standard (UAS) as a 27-requirement regulatory profile extension, comprising 14 system, 8 human, and 5 bridge requirements, without altering the core 161-control structure or creating CP.11.

— Gives exact requirement counts and scope boundaries, no usage steps

2026-09-04_Agent_CLI_AISM_Decision_Supp…

THINNER COVERAGE BELOW

03 **Evidence source integrations for NEXUS and NVIDIA SkillSpector**

NEW

55

Adds NEXUS integration so NEXUS findings can be incorporated as structured evidence within AI SAFE² assessments, connecting implementation observations to governance and assurance requirements. Also adds NVIDIA SkillSpector as an evidence source for

assessing agent skills and implementation risks, with source attribution and evidence context retained.

— Names both integrations and their evidence roles, no config steps given

2026-09-04_Agent_CLI_AISM_Decision_Supp…

04 **AISM Decision Cards for assessment findings** NEW 45

Adds AISM Decision Cards that transform technical assessment findings into structured decision-quality summaries covering scores, evidence quality, conflicts, alternatives, pros/cons, and recommended paths forward.

— Describes content of cards but no interface or command to invoke

2026-09-04_Agent_CLI_AISM_Decision_Supp…

WAS THIS USEFUL? 👍 👎

Aurascape

Feedback ⓘ 1 RELEASE · 2026-09-04 BLOG ↗ RANK

Aurascape is the AI-native security platform that empowers enterprises with real-time visibility, intent-based controls, and protection for every AI interaction—securing innovation without slowing it down.

> Aurascape launched Coding Assistant Guardrails, a new solution that inspects and enforces policy on live AI coding sessions across Claude Code, GitHub Copilot, and Cursor, blocking secrets, destructive commands, and prompt injection in real time before any artifact exists to scan.

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

01 Destructive command and unauthorized git push interception

NEW 80

Intercepts destructive shell commands and unauthorized git pushes at the wire level, in the tool call the model streams back, before the client executes them — with configurable block or confirmation-checkpoint policy.

— Describes wire-level mechanism and configurable policy options

Secure the Coding Session, Not Just the…

02 Inline decoding of coding-assistant streaming protocols NEW

75

Decodes coding-assistant streaming protocols — Server-Sent Events, WebSockets, Protobuf, gRPC, and MCP — inline without buffering, preserving native streaming speed for developers while enabling content inspection.

— Names every protocol and explains the no-buffering mechanism

Secure the Coding Session, Not Just the…

03 Continuous discovery of IDE plugins and MCP servers NEW

70

Provides continuous discovery of IDE plugins, extensions, and MCP servers as they appear, with new or uncataloged apps receiving signatures on a 48-hour SLA — removing the need for hand-maintained allowlists.

— Names the SLA and removes-allowlist mechanism concretely

Secure the Coding Session, Not Just the…

04 Coding Assistant Guardrails for live sessions NEW

65

Aurascape launches Coding Assistant Guardrails, a new solution that inspects and enforces policy on live AI coding-assistant sessions — covering Claude Code, GitHub

Copilot, and Cursor — before any artifact exists to scan.

— Names the covered tools but not the enforcement mechanism

Secure the Coding Session, Not Just the…

05 **In-flight secrets detection in prompts** NEW 65

Detects and blocks or redacts secrets and credentials pasted into prompts in flight, before the model receives them, with the event captured as auditable evidence.

— Clear mechanism and timing, no config surface named

Secure the Coding Session, Not Just the…

06 **Auditable conversation logs with RBAC** NEW 60

Delivers auditable conversation logs — capturing prompts, responses, connections, and tool calls behind each policy flag — under role-based access control for security investigations.

— Names what's captured and the RBAC control, no UI path

Secure the Coding Session, Not Just the…

THINNER COVERAGE BELOW

07 **Prompt injection detection in docs and MCP outputs** NEW 55

Detects prompt injection hidden in READMEs, documentation, and MCP outputs, neutralizing hijacked instructions before the agent acts on them.

— Names sources of injection but not detection mechanism

Secure the Coding Session, Not Just the…

08 **Customizable in-context coaching templates** NEW 50

Offers customizable coaching templates that surface in-browser, in-IDE, or in-terminal explanations and safer paths when a policy action fires.

— Names surfaces where coaching appears, no configuration detail

Secure the Coding Session, Not Just the…

09 **Typosquatted package and spoofed link flagging** NEW 45

Flags typosquatted packages and spoofed links recommended by a model, surfacing supply-chain attempts as visible, reportable events.

— Bare description of what is flagged, no mechanism given

Secure the Coding Session, Not Just the…

10 **MCP tool call parameter visibility** NEW 40

Exposes MCP tool calls down to their parameters for visibility and policy enforcement.

— Single-line description with no mechanism or scope

Secure the Coding Session, Not Just the…

11 **Natural-language queries via Auri** NEW 35

Enables natural-language queries over AI usage, risk, and policy data through Auri.

— Single sentence with no example query or mechanism

Secure the Coding Session, Not Just the…

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ 1 RELEASE · 2026-09-05 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian 1.14.0 brings native iOS Quick Capture with Live Activity support, alongside a kanban layout in Bases, editor color-highlight swatches, and RTL workspace support.

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

01 **iOS Quick Capture with Live Activity** NEW 75

Adds native Quick Capture on iOS 26, allowing note capture from the Lock Screen, Control Center, or Shortcuts without loading the vault; active captures appear as a Live Activity on the Lock Screen and Dynamic Island. Quick Capture templates also gain `{{latitude}}` and `{{longitude}}` variables for inserting the current GPS position into captured notes.

Tag a field-note highlight with a location when capturing outdoors, so coordinates are embedded automatically via Quick Capture.

> 📍 In Quick Capture, create a Location pointing to today's daily note and set the template body to: ==🔴 GPS fix at {{latitude}}, {{longitude}}==

– Names the OS version, entry points, and template variables with an example.

changelog-20260905-05a332ef

THINNER COVERAGE BELOW

02 **Per-line CSS targeting in math blocks** IMPROVED 55

Each line of a math block (`$$`) now carries the `.HyperMD-math` CSS class, enabling per-line targeting by themes and plugins.

– Names the exact CSS class and syntax but no usage example. changelog-20260905-05a332ef

03 **Color swatch in Live Preview highlights** NEW 50

Adds an inline color swatch in Live Preview that appears when the cursor overlaps a highlight, allowing quick color changes.

– Describes behavior but only via UI, no example given. changelog-20260905-05a332ef

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 5, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue