THE CYBER TOOLCHAIN

NO. 099

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 9, 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 14 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 8 tools

▾ // THE BRIEF HIDE

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. GitGuardian can now scan Harbor registries for leaked secrets and answer, on demand, whether a given is already a known leak. Auth0 made per-request scope trimming

vim ? help

in Post Login Actions generally available, and closed a legacy path that let third-party apps widen scope past consent. Tines gave AI clients write access to Records.

PROTECT

Find secrets baked into container images, and check a suspect string without filing a ticket

Harbor registry scanning extends secret detection past source control into built images, where credentials get copied in during the build and nobody looks again. The on-demand presence-check API lets an incident responder or a pipeline ask whether a specific secret is already known-leaked, instead of waiting for a scan cycle; per-secret environment classification in NHI Governance separates a prod key from a sandbox one so triage stops treating them alike.

GitGuardian SaaS

PROTECT

Mint access tokens with only the scopes the request needs

Scope customization reaching GA in Post Login Actions means the token issued at login can be trimmed per user, app or context in code rather than relying on static API permission grids, and Actions can now read the live scope state before deciding. The restriction on the legacy scope API for third-party apps shuts a route where an app could end up with scopes the user never consented to.

Auth0

RESPOND

Let an AI client write back to Tines Records, not just read them

Records are Tines' persistent key-value store, commonly used for case state, suppression lists and enrichment caches. MCP write support means an agent handling a triage loop can update that state directly — mark a case, add an indicator, create a record type — instead of a human pasting results into a workflow run.

Tines

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ

Scan your Harbor container registry for secrets (2026-10-09) BLOG ↗ RANK

GitGuardian SaaS is a cloud-based security platform that detects and prevents secrets leakage in git repositories and source code.

> GitGuardian SaaS adds Harbor registry scanning, an on-demand secret presence-check API, a new is all of source-status filter, and per-secret environment classification in NHI Governance, alongside PR comment notes and Google Cloud keyless authentication.

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

01 Harbor registry scanning integration NEW INTEGRATION NEW

90

GitGuardian can now connect to self-hosted Harbor registries to scan image layers for secrets. Setup is via Integrations > Harbor: create a system-level robot account (with the `robot$` prefix) with Repository (List, Pull) and Artifact (List, Read) permissions, enter the Harbor URL and credentials, then choose to monitor specific repositories or the entire instance.

Set up Harbor integration via the GitGuardian dashboard to begin scanning all image layers in a self-hosted registry for secrets.

> 📍 1. In Harbor, create a system-level robot account with Repository (List, Pull) and Artifact (List, Read) permissions. 2. In GitGuardian, go to Integrations › Harbor. 3. Enter your Harbor URL and the robot account name (including its `robot$` prefix) and secret. 4. Choose your monitored perimeter: select specific Harbor repositories or monitor the entire instance.

— Full setup steps and required permissions given product docs

02 On-demand secret presence check API NEW

90

Adds `POST /v1/incidents/secrets/{id}/check-presence` and `POST /v1/public-incidents/secrets/{id}/check-presence` endpoints to trigger an asynchronous presence check on a secret incident, letting you verify whether a leaked secret is still live without waiting for the next scheduled scan.

Trigger an on-demand presence check for a specific secret incident without waiting for the next scheduled scan.

```
$ curl -X POST
'https://api.gitguardian.com/v1/incidents/secrets/42/check-
presence' \
  -H 'Authorization: Token <your-api-token>'
```

— Exact endpoints plus a runnable curl example product docs

03 **Source status 'is all of' filter** NEW 85

Adds `source_monitoring_status__all` operator to the public API, exposing an `is all of` filter so you can query incidents whose sources are ALL deleted or archived, e.g. `source_monitoring_status__all=deleted,archived`. The same `is all of` operator is now also available on the `source status` filter in the UI.

Filter the incidents API to only surface incidents where ALL sources are deleted or archived, reducing noise from partially-removed code.

```
$ curl -G 'https://api.gitguardian.com/v1/incidents' \
  -H 'Authorization: Token <your-api-token>' \
  --data-urlencode
'source_monitoring_status__all=deleted,archived'
```

— Exact query parameter and runnable example given product docs

04 **@gitguardian PR comment notes** NEW 75

Developers can reply to a GitGuardian check run by commenting `@gitguardian` on a pull request; the comment is added as a note on every incident the check run found, visible in the dashboard and through the incident notes API.

— Exact trigger phrase named, no API path given product docs

05 **Per-secret environment classification in NHI Governance**

IMPROVED 70

NHI Governance now reads environment classification per secret from secret paths, tags, and labels, recognising `prod`, `prd`, `stg`, `qa`, `uat`, `dev`, and more, before falling back to the source-level `env` setting when no per-secret signal is found.

— Recognized values named but no usage path given product docs

THINNER COVERAGE BELOW

06 **Google Cloud keyless authentication** NEW 55

Supports connecting GitGuardian to Google Cloud keyless via Workload Identity Federation, removing the need to store or rotate a service account key, with project,

folder, or organization-level scopes.

— Mechanism and scopes named but no setup steps product docs

07 **Confluence notifications use stable user key** IMPROVED

40

Confluence incident notifications now carry the stable Confluence user key for new occurrences, enabling mapping of incident authors to users even when the instance does not expose emails.

— Thin description, no config or action named product docs

08 **GitLab System Hooks non-admin token coverage** IMPROVED

40

GitLab System Hooks integrations with a non-admin token now monitor projects in personal namespaces and in groups the service account reaches only through a project.

— Thin description, no config key or steps named product docs

└──▷ ALSO FROM THESE RELEASES

Scan your Harbor container registry for secrets (2026-10-09) BLOG ↗ LATEST

STABLE

Harbor integration uses a read-only robot account, supports air-gapped instances via GitGuardian Bridge, flags unhealthy installs, and enables secret rotation without reinstall.

Uses a Harbor robot account with Repository (List, Pull) and Artifact (List, Read) permissions only — read-only, scoped, and revocable at any time.

Supports scanning Harbor instances not reachable from the internet via GitGuardian Bridge.

Allows selecting specific Harbor repositories or monitoring the entire instance with automatic inclusion of new repositories.

Performs both historical scanning of existing images and incremental scanning of newly pushed images.

Flags unhealthy installations (expired robot accounts, lost permissions, or unreachable instances) and supports rotating the robot account secret without reinstalling.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Endor Labs Feedback ⓘ

Endor Labs + SpaceXAI: Securing every stage of agentic software delivery | Blog | Endor

BLOG ↗

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC.

Endor Labs Agent Kit brings AURI remediation agents (AI SAST and SCA) into Grok Build sessions and CI, with per-change approval gates and audit-ready evidence.

NEW INTEGRATION Adds Endor Labs Agent Kit for Grok Build, installable via `/add-plugin endorlabs` in the agent chat or from the marketplace, bringing specialized AURI remediation agents directly into the IDE and CI/cloud runs through the Grok Build SDK.

Adds AI SAST Remediation agent (part of the Endor Labs Agent Kit) that traces data flow from untrusted input to vulnerable calls, proposes fixes and test plans, and can open a pull request upon developer approval.

NEW DETECTION Adds SCA Remediation agent (part of the Endor Labs Agent Kit) that queries scan results, checks package usage, applies upgrades or Endor Patches, adds tests, and rescans to confirm no new vulnerabilities were introduced.

NEW DETECTION Adds Endor Scores rating system for MCP servers and skills, mapped to the OWASP Top 10 for LLM Applications, giving security teams a trust signal for every agent component in use.

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ Add AI Runtime Control product (2026-09-22) BLOG ↗

RANK

Fastly's edge cloud platform delivers faster, safer, and more scalable sites and apps to customers. Elevate your edge CDN, video delivery, security, and more.

> Fastly's product enablement API now supports enabling and disabling the new AI Runtime Control product programmatically.

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

01 **AI Runtime Control product identifier** NEW INTEGRATION NEW

55

The Fastly product enablement API adds `ai_runtime_control` as a supported product identifier, enabling programmatic enable/disable of the AI Runtime Control product.

— Names exact API identifier but no endpoint or usage detail

Add AI Runtime Control product (2026-09…

Route, control, and monitor AI model traffic with the new AI Runtime Control API (2026

BLOG ↗ LATEST STABLE

New AI Runtime Control API gateway adds virtual key auth, per-key rate limits, failover policies, usage metrics with CSV export, and session logs.

Adds the AI Runtime Control API gateway — point existing SDKs at Fastly by changing only the base URL and credentials to gain per-key controls and usage reporting without rewriting application code.

Adds provider connection management to supply credentials (API key or, for Amazon Bedrock, an AWS IAM role) that AI Runtime Control uses to reach a provider on your behalf; also exposes endpoints to look up supported providers and models and retrieve Bedrock IAM setup information.

Adds virtual key management with per-key `request-per-minute` and `token-per-minute` limits to control how each application authenticates to and consumes the gateway.

Adds failover policy configuration on a virtual key to automatically retry against alternate provider connections after timeouts, rate limiting, or server errors.

Adds usage metrics for requests, sessions, and token consumption, filterable by virtual key, provider, or model, with CSV export.

+1 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SOAR

Tines Feedback ⓘ

Create, update, and delete Records via the Tines MCP (2026-10-09) BLOG ↗

RANK

Tines lets teams put AI to work across their most sensitive systems, data, and processes, with the security and governance to do it safely.

> Tines extended its MCP server to support full write operations on Records, letting AI clients create, update, and delete records and record types rather than just read them.

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

01 **Write operations for Records via Tines MCP** IMPROVED

50

The Tines MCP server now supports creating, updating, and deleting records and record types, extending it beyond read-only access so AI clients can perform full write operations.

— Describes what changed but no mechanism, endpoint or limits given.

Create, update, and delete Records via …

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir Feedback ⓘ

Tenzir Node v6.20.0-rc.1: A new data model for faster pipelines (2026-10-09)

BLOG ↗ RANK

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

> Tenzir Platform v1.40.2 updates the Explorer to correctly group events under the new data model.

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

01 **Explorer grouping by schema name** NEW DETECTION IMPROVED

40

The Explorer now groups events by schema name and displays accurate per-name counts, even when field sets differ across events, reflecting the new data model.

— Describes behaviour but no UI path or config detail

Tenzir Platform v1.40.2: Explorer suppo…

UPDATED INTEGRATION SentinelOne queries through the console API

UPDATED INTEGRATION Filter and time-window pushdown for SentinelOne

NEW INTEGRATION Windows Event Forwarding without a Windows collector

NEW INTEGRATION Read from and write to DuckDB databases

UPDATED INTEGRATION Filter, projection, and limit pushdown for MySQL and SQL Server

> Tenzir Node v6.20.0-rc.1: A new data model for faster pipelines (2026-10-09)
>
> BLOG ↗ PRE-RELEASE
>
> New heterogeneous-list data model, typed `ai_decide` decisions, encryption functions, Luhn/CRC-32 checksums, and secrets as function arguments.
>
> Adds typed decisions with `ai_decide`, enabling structured, type-safe outputs from AI decision steps in pipelines.
>
> Adds encryption functions to the pipeline expression language.
>
> Adds Luhn check digit and CRC-32 checksum functions.
>
> Adds `hmac` secret keys and Crypto-PAn secret seeds as function arguments, expanding secrets into cryptographic functions.

> Supports secrets as function arguments, allowing secret values to be passed directly into pipeline functions.
>
> +6 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Zama fhEVM

Feedback ⓘ v0.14.3 NOTES ↗ RANK

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

> fhEVM v0.14.3 made the e2e test suite's coprocessor configuration generation driven by environment variables, easing CI setup.

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

01 **Env-var driven e2e coprocessor config generation** IMPROVED

35

The test suite now generates e2e coprocessor configuration from environment variables, enabling automated and CI-friendly stack setup.

— No specific env var names or commands given v0.14.3

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Auth0 Feedback ⓘ

Actions - Access Token Scope Customization - GA (2026-10-09) BLOG ↗ RANK

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

> Auth0 promoted Access Token Scope Customization in Actions to General Availability, extending it to Post Login Actions and exposing live scope state, while restricting a legacy scope API for third-party applications to prevent consent bypass.

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

01 **Access Token Scope Customization GA with Post Login support**

IMPROVED 75

Access Token Scope Customization for Actions — including the Add, Remove, Set, and Clear scope interfaces — reached General Availability for Credentials Exchange Actions and was extended to Post Login Actions. Actions can now read `event.transaction.target_scopes` to see the live target scope list, which reflects every change made across chained Actions.

.
.
.
.

— Names the exact scope interfaces and the event property, but no code sample.

Actions - Access Token Scope Customizat…

02 **Third-party addScope() restricted to prevent consent bypass**

70

BREAKING

api.accessToken.addScope() is no longer available in Actions executed for third-party applications. Any existing Action relying on it for third-party apps must be migrated to the target scope interfaces (Add, Remove, Set, Clear).

— Names the exact API call and required migration path, no code sample.

Actions - Access Token Scope Customizat…

└──▷ BREAKING ON UPGRADE

! api.accessToken.addScope() is no longer available in Actions executed for third-party applications; any existing Action relying on it for third-party apps will need to be migrated to the target scope interfaces.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Sep 23rd, 2026 NOTES ↗

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

Overnight integration syncs now stagger their start times to reduce database load during high-volume windows.

UPDATED INTEGRATION AWS integration — six new evidence collection areas

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 9, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue