THE CYBER TOOLCHAIN

NO. 126

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 31, 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 114 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 36 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.

PROTECT

Rip out the eMudhra intermediates and find out who was trusting m

vim ? help

Codezero's Cordon adds a CA bundle refresh command and hardens its TLS verifier so the compromised intermediates are no longer accepted, which is the part most teams would otherwise hand-script across hosts. Cycode's Threat Intelligence dashboard queries turn the same incident into a list of affected assets tied back to their owning repository, so the containment work has an owner instead of a mailing-list announcement.

Codezero · Cycode

DETECT

Keep kernel tracing policies loaded across restarts and upgrades

Tetragon v1.8.0 adds stateful and persistent tracing policies, so a policy no longer has to be re-applied after the agent cycles — historically the gap where an attacker gets a free window during a node reboot or operator upgrade. mTLS on the gRPC server and uprobe binary digest verification also mean the event stream and the probed binaries can be authenticated rather than assumed.

tetragon

PROTECT

Make SMS MFA abuse expensive without blocking real users

SMS Fraud Protection swaps static rate limits for exponential backoff on Phone MFA, which is the right shape for toll fraud and SMS pumping: an attacker hammering a number hits rising delay while a legitimate retry barely notices. Throttled attempts land in audit logs, so you can see the attack rather than only its carrier bill.

Auth0

IDENTIFY

Prove which tenant and account a compliance scan actually ran as

ScubaGear's new command confirms the account and tenant behind a scan, which sounds trivial until a multi-tenant M365 assessment produces a clean report against the wrong directory. Removes the manual credential-context check before you hand results to an auditor.

ScubaGear

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

julius

Feedback ⓘ v1.5.5 NOTES ↗ RANK

Simple LLM service identification - translate IP:Port to Ollama, vLLM, LiteLLM, or 60+ other AI services in seconds

> julius's probe command now flags LLM endpoints that require authentication, helping practitioners quickly separate protected services from open ones during scanning.

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

01 **Auth-required signal in probe fingerprint output** NEW

45

The `probe` command now includes an auth-required signal in its fingerprint output, letting practitioners immediately distinguish protected LLM endpoints from open ones during scanning.

— Names the command and signal but no output format or example v1.5.5

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

changelog · Oct 31, 2026 CHANGELOG ↗ RANK

GitGuardian Detection Engine scans source code repositories, commits, and pipelines to detect leaked secrets such as API keys, passwords, and credentials.

> GitGuardian Detection Engine tightened checker requirements and refined two individual secret detectors in this release.

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

01 **Custom host requirement for eight checkers** BREAKING

55

Eight checkers now need a custom host to run, including JFrog Artifactory, AWS Cognito, Bitbucket, Jira, Azure Storage, and Okta.

— Names six of eight affected checkers but no config detail changelog-20261031-214c3208

02 **Jenkins API token detector matches instance URL** IMPROVED

55

The `jenkins_api_token` detector now also matches the Jenkins instance URL, in addition to the token itself.

— Clear mechanism change, no further scope given changelog-20261031-214c3208

03 **Azure storage key detector split into own group** IMPROVED

45

The `microsoft_azure_storage_account_key_with_account_name` detector moves to its own detector group.

— Names detector but not the new group's purpose changelog-20261031-214c3208

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Veracode

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Veracode is a cloud-based application security platform that scans code, binaries, and running applications to find and help remediate software vulnerabilities.

> Veracode's October 2026 release adds asynchronous Fire-and-Forget remediation to its CLI, expands DAST authentication and compliance evidence collection, and broadens Static Analysis Engine language support to JDK 27, Spring AI, and Scala 3.9.

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

01 **Fire-and-Forget async mode in Veracode CLI** NEW 75

Veracode CLI v2.54.0 adds asynchronous Fire-and-Forget (FnF) mode for the `fix` command, submitting remediation requests for background processing without waiting for fixes to be generated and pull requests to be created.

– Names CLI version and command but no flag syntax shown. changelog-20261031-c1c7a4a2

THINNER COVERAGE BELOW

02 **Domain field for DAST HTTP basic auth** NEW 55

DAST web application scans now support an optional Domain field when configuring HTTP basic authentication, enabling authentication to applications that require domain-based credentials.

– Names exact config field but not full setup path. changelog-20261031-c1c7a4a2

03 **Compliance evidence collection for GDPR, NIS2, DORA** NEW 55

Scans now collect compliance evidence covering cookie consent tracking, third-party data flows, destination countries, and privacy, legal, and copyright documentation, supporting GDPR, NIS2, and DORA assessments.

– Describes scope but no config or endpoint to act on. changelog-20261031-c1c7a4a2

04 **Static Analysis Engine v2026.10 language and framework support**

NEW 45

Static Analysis Engine v2026.10 adds support for `JDK 27`, LLM-powered applications built with Spring AI, and `Scala 3.9`.

– Lists supported versions but no scan configuration detail. changelog-20261031-c1c7a4a2

WAS THIS USEFUL?  

◆ Software Supply Chain Security

Cycode

Feedback ⓘ

node2shell / HTTP2Shell (CVE-2026-111497): What to patch and how to check if you're expos

BLOG ↗ RANK

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

> Cycode shipped two rounds of Threat Intelligence dashboard queries to help teams find exposure to emerging threats: CVE-2026-111497 (node2shell/HTTP2Shell) across Node.js, Bun, .NET and Python runtimes, and the eMudhra intermediate CA compromise, with matched assets linked to their owning repository via the Context Intelligence Graph.

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

01 **eMudhra CA compromise detection and manual remediation guidance**

NEW 93

Three new Context Intelligence Graph queries in the Threat Intelligence dashboard find container images with unpatched `ca-certificates`, images with outdated runtimes (Node.js, Java, pip, certifi), and Dockerfiles on unpatched base images — each showing matched asset counts and linking to the originating repository. Runtimes can be tested for continued trust in the compromised eMudhra intermediates by connecting from inside containers or CI runners to `https://dv-g3b-emudhra.we-have-a-problem.com` and `https://dv-g3b-aaa-emudhra.we-have-a-problem.com` via `curl`, Python's `urllib.request`, or Node.js `fetch` (a certificate error means protected, a page load means still exposed). Custom CA bundles set via `SSL_CERT_FILE`, `REQUESTS_CA_BUNDLE`, `CURL_CA_BUNDLE`, `PIP_CERT`, or `-Djavax.net.ssl.trustStore` are never updated by the OS and require manual removal of the 16 compromised eMudhra intermediates.

Verify from inside a container or CI runner that the curl/OS trust store no longer accepts certificates chaining to the compromised eMudhra intermediates — a certificate error is the passing result.

```
$ curl -sS https://dv-g3b-emudhra.we-have-a-problem.com
```

Check whether the Python runtime in your environment or image still trusts the compromised CAs via its OS trust store.

```
$ python -c "import urllib.request;
urllib.request.urlopen('https://dv-g3b-emudhra.we-have-a-
problem.com'); print('STILL TRUSTED')"
```

Rebuild a container image so that the base image and ca-certificates package layers refresh to a patched version.

```
$ docker build --pull --no-cache .
```

— Includes runnable curl/Python test commands and exact query surfaces

eMudhra intermediate CA compromise: How…

02

CVE-2026-111497 detection queries in Threat Intelligence dashboard

NEW

83

Five new queries were added to the Threat Intelligence dashboard for node2shell/HTTP2Shell (CVE-2026-111497): a 'services that accept Brotli bodies' query finding Dockerfiles on a vulnerable Node.js or Bun image that also have, in the same directory, a dependency that decompresses Brotli request bodies (Express 5, body-parser 2, Koa body parsers, or `@fastify/compress`); a Dockerfiles on vulnerable runtimes query finding Dockerfiles building on Node.js below 22.23.4, 24.22.1, or 26.12.1, end-of-life Node.js lines, or Bun below 1.4.4, including digest-pinned images and `nvm install` or NodeSource setup lines; two container image queries finding built images whose Node.js/Bun runtime is below the fixed version or which ship the vulnerable distribution Brotli library; and a code dependencies query finding projects (including indirect dependencies) on vulnerable Brotli decoder libraries — `brotli`, `brotlicffi`, and `curl_cffi` in Python, `com.aayushatharva.brotli4j` in Java, and the Ruby `brotli` gem. Each dashboard item links matched assets to their owning repository through the Context Intelligence Graph.

— Names exact versions, libraries and queries but no runnable command

node2shell / HTTP2Shell (CVE-2026-11149…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Abnormal AI

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗

Abnormal AI is a behavioral AI security platform that protects email, identity, and AI—and stops insider threats. Trusted by 5,000 organizations.

Abnormal AI adds time-series identity posture tracking, per-signal reset risk detail, AI Data Analyst scheduling with a Report Repository, and Segmented SaaS multi-tenant isolation.

Identity Threat Protection now shows security grade and threat exposure over time, letting admins review up to 180 days of history with markers for improvements, regressions, and coverage changes, and filter identities at risk by attack scenario such as Helpdesk Impersonation, Vendor Email Compromise, and MFA Fatigue Push.

Identity Threat Protection now exposes the individual risk signals behind each Adaptive Identity Reset risk score, with per-signal ratings on a four-tier scale (Low, Moderate, Elevated, or High) and plain-language definitions alongside the identity quiz result.

AI Data Analyst adds scheduled reporting — one-time or recurring — with a Schedule as Custom Report option that derives the reporting window from a chosen frequency, plus a new Report Repository for versioned downloads and full conversation history.

AI Data Analyst adds a single multi-tenant management view for MSSPs and multi-tenant customers to manage reporting across every tenant from one place.

Abnormal adds Segmented SaaS, letting organizations create separate Abnormal tenants for business units, subsidiaries, regions, or agencies that share one email tenant, with users assigned by email domain, directory group, or a list of user principal names, and each tenant getting its own admins, views, alerts, and data over a single integration.

+1 more

UPDATED DETECTION Malicious link detection — URL verdict scoring

NEW DETECTION Visual detection model for image-based brand impersonation

UPDATED DETECTION Malicious link detection — URL-hosting platform remediation (Week of Oct 30)

UPDATED DETECTION Brand impersonation detection — visual deception evasion

NEW DETECTION Sender Compromise model

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Plexicus

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Plexicus validates security findings with evidence and helps your team fix what matters. AI Swarm Pentest, Deep Code Analysis, Remediation and XBOM.

> Plexicus reworked its AI Pentest experience with a live War Room and per-node automation testing, alongside a new Scan Modules settings page, a granular role permission matrix, a unified compliance view, and richer certification question types.

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

01 AI Pentest wizard with live War Room and asset views IMPROVED

70

The AI Pentest wizard was reworked with a live war-room board for real-time visibility during scans, per-node automation testing with data mapping, and repo-centric asset views that surface repository-level context during pentests.

— Mechanism and scope described but no exact UI path or API given

changelog-20261031-41e91c48

THINNER COVERAGE BELOW

02 Scan Modules settings page replaces connector catalog NEW

50

Plexicus now has a Scan Modules settings page that replaces the old connector catalog.

— Names the replaced surface but no further mechanism detail changelog-20261031-41e91c48

03 Role permission matrix for platform access NEW

45

A View/Edit/Manage role permission matrix now controls per-role access across the platform.

— Names permission levels but no navigation path or config key

changelog-20261031-41e91c48

04 Richer certification question types and lifecycle IMPROVED

25

Plexicus now supports richer question types and lifecycle management for certifications.

— States the addition but not which question types or lifecycle stages

changelog-20261031-41e91c48

05 **Unified compliance controls view** IMPROVED 20

Plexicus now consolidates compliance controls into a single unified view.

— Bare description with no mechanism or location specified changelog-20261031-41e91c48

WAS THIS USEFUL? 👍 👎

Lacework FortiCNAPP

Feedback ⓘ updated · Oct 31, 2026

NOTES ↗ RANK

Lacework FortiCNAPP is a cloud-native application protection platform from Fortinet that secures cloud workloads, containers, and infrastructure through vulnerability scanning, misconfiguration detection, and runtime threat monitoring.

> Lacework FortiCNAPP's October update adds context to intrusion graphs and refines its compliance evaluation engine.

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

01 **Connection-reason visibility in intrusion graphs** IMPROVED 45

The intrusion graph now shows why entities in a composite alert are connected, giving analysts direct context for relationships between flagged resources.

— Explains what changed but no UI path or mechanism given snapshot-20261031

02 **Improved compliance evaluation engine** IMPROVED 7

Lacework FortiCNAPP shipped an improved compliance evaluation engine, though no further detail on the change is provided.

— Mentioned only in summary with no detail or mechanism snapshot-20261031

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Cilium tetragon

Feedback ⓘ v1.8.0 NOTES ↗ RANK

Tetragon is an eBPF-based security observability and runtime enforcement tool for Kubernetes and cloud-native environments.

> Tetragon v1.8.0 introduces stateful and persistent tracing policies, mTLS for the gRPC server, and uprobe binary digest verification, alongside a wave of new tracing-policy selectors, CEL expression operators, uprobe actions, and BPF sensor configuration options, plus several breaking changes to deprecated fields, actions, and metrics labels.

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

01 **Tracing policy selectors and infrastructure** NEW 80

Adds `matchCmdArgs`, `matchWorkloads`, `nodeSelector`, and `MatchCEL` tracing policy selectors for filtering by command-line arguments, workload, Kubernetes node, and mixed CEL expressions, alongside new stateful tracing policies, gRPC-based persistent policies, JSON schema and CEL validation for tracing policies without Kubernetes support, an eBPF-based glob library for tracing policy matching, and domain-sharded tracing policy management.

— Four selectors named but no config example shown Mechanism named for five additions but no commands shown

v1.8.0

02 **BPF sensor configuration and argument types** NEW 80

BPF sensors gain `btfType` and `btfTypeModule` argument types, support for the `fd` type to retrieve `file*`, the `--sleepable-preload-size` option to configure the sleepable preload size, a read-only frozen configuration map, the `--execve-map-no-prealloc` option to control execve map memory allocation, BPF area-specific DEBUG() rodata configuration, and support for multiple preloaded string-resolved arguments in BPF sensors.

Prevent execve map from pre-allocating memory at startup to reduce idle memory usage on nodes with many short-lived processes.

```
$ tetragon --execve-map-no-prealloc
```

— Concrete flags named, one has a runnable example v1.8.0

03 **Deprecated tracing policy syntax removed** BREAKING 70

The deprecated `sizeArgIndex` field is replaced by `sizeFromArgIndex` in tracing policies, and policies using the old field must be updated. Previously deprecated actions have been removed entirely and policies referencing them are now rejected on upgrade.

— Exact field names given plus upgrade impact `v1.8.0`

---

04 **New uprobe actions and verification** NEW

65

Uprobes gain a `Set` action with `u64 argValue` support, `ProcessUprobe` events now include symbolized user-space stack traces for uprobe Post actions that enable `userStackTrace`, and Tetragon adds uprobe binary digest verification.

— Named action and field but no usage example Only named in summary, no mechanism described

`v1.8.0`

---

05 **CEL expression operators and register overrides** IMPROVED

65

CEL expressions in tracing policies gain support for register overrides via `argRegs`, bit-shifting operators (`lsh` , `rsh`), 16-bit and 8-bit signed and unsigned integers, multiplication and division operators, and bitwise operators (`and` , `or` , `xor` , `not`).

— Six CEL additions named without example expressions `v1.8.0`

---

06 **Postfix/NotPostfix length limit on kernels without large BPF program support**

65

BREAKING

`matchArgs` / `matchBinaries` Postfix and NotPostfix values longer than 87 characters are now rejected on kernels without large BPF program support.

— Exact limit and condition given, clear impact `v1.8.0`

---

07 **Metrics label changes for opcode and event type** BREAKING

65

Metrics labels changed: `opcode` now carries the numeric `ops.OpCode` value, `opstr` is the new label for its readable name, and `event_type` replaces the prior `EventType` labeling; dashboards or alerts keyed on the old label names will break.

— Exact label names given but no migration steps `v1.8.0`

---

08 **Uprobe selector and override action improvements** IMPROVED

60

Tracing policies can now filter uprobe selectors by the caller function in addition to the called function, and the uprobe override action gains `argNew{Symbol,Addr,Offset}` support plus support for multiple `sopath` actions.

— Thin one-line addition with no syntax or example Named fields but no example of override syntax

v1.8.0

THINNER COVERAGE BELOW

09 **gops debug interface disabled by default** BREAKING 40

The gops debug interface is now disabled by default in the Helm chart.

— Single line, no reinstatement instructions given v1.8.0

10 **mTLS for gRPC server** NEW 25

Tetragon's gRPC server now supports mTLS.

— Only named in summary, no configuration details given v1.8.0

└──▷ **BREAKING ON UPGRADE**

! The deprecated `sizeArgIndex` field is replaced by `sizeFromArgIndex`; policies using `sizeArgIndex` must be updated.

! Deprecated actions have been removed; policies referencing them will be rejected on upgrade.

! `matchArgs` / `matchBinaries` Postfix and NotPostfix values longer than 87 characters are now rejected on kernels without large BPF program support.

! The gops debug interface is now disabled by default in the Helm chart.

! Metrics labels are changed: `opcode` now carries the numeric `ops.OpCode`, `opstr` is the new label for its readable name, and `event_type` replaces prior EventType labeling; dashboards or alerts keyed on the old label names will break.

WAS THIS USEFUL?  

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Protection for Linux

Feedback ⓘ

Version 5.22.0New NOTES ↗

Sophos Protection for Linux is an endpoint security product that detects and blocks malware, ransomware, and suspicious behavior on Linux servers and workstations.

Sophos Protection for Linux adds Helm options to exclude selected processes from sensor events and selected paths from file events.

Exclude selected processes from process, connection, and DNS events by configuring a new Helm option.

Exclude selected paths from file events by configuring a new Helm option.

NEW SUPPORTED PLATFORM Kernel 7.3 support

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.261030.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI's latest release adds dashboard variable restriction modes, Hub-based custom view import, configuration bundle export/import, manual security coverage updates, a pingConnector health-reporting method in pycti, full timestamp display in tables, EE-only custom-field management, and an in-progress ingestion health monitoring system behind a feature flag.

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

01 **pingConnector for connector health in pycti** NEW 70

The `pingConnector` method in pycti lets connectors report errors back to OpenCTI, improving connector health visibility.

— Names exact method, usable directly in pycti code 7.261030.0

02 **Ingestion health monitoring behind feature flag** NEW 60

Prepared ingestion health monitoring including event type and scope filter values, email notifications, guards, service account checks, and a UI, currently gated behind a feature flag.

— Lists components but feature is not yet enabled 7.261030.0

THINNER COVERAGE BELOW

03 **Import custom views from Hub** NEW 50

Custom views can now be imported directly from the Hub using the new `import-from-Hub` button in custom views.

— Names exact UI button but no further mechanism 7.261030.0

04 **Configuration bundle export and import** NEW 50

Themes can now be exported as part of configuration presets, allowing platform appearance settings to be packaged and shared; a configuration bundle can also be imported from the bootstrap wizard, enabling faster initial platform setup.

— Describes two bundle actions but no format or key names 7.261030.0

05 **Manual security coverage result updates** NEW 45

Analysts can now manually update security coverage results, giving direct control over coverage data without waiting for automated runs.

— Explains benefit but no UI path or field names 7.261030.0

06 **Full timestamp display in table views** IMPROVED 45

Table views can now show full timestamps (date and time) instead of date-only, with configurable display.

— Mentions configurability but not where to set it 7.261030.0

07 **Custom-field definition management restricted to EE** BREAKING 40

Management of custom-field definitions is now restricted to Enterprise Edition instances.

— States the restriction without migration guidance 7.261030.0

08 **Variable restriction modes for dashboards** NEW 28

Added restriction modes to control how variables can be used within dashboards.

— No detail on what the modes are or how they work 7.261030.0

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Claude Code auto mode for Anthropic and non-Anthropic models (2026-10-30)

BLOG ↗ RANK

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, machine learning, and AI workloads, built on Apache Spark.

> Databricks shipped GA real-time mode for Lakeflow pipelines with sub-5ms latency, GA native XML ingestion and parsing, beta relationship constraints between Unity Catalog metric views, new audit logging and dashboard visibility for Unity Gateway Skill usage, and Claude Code auto mode support for Anthropic and non-Anthropic models.

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

01 **Relationships between metric views (Beta)** NEW 90

You can now define how Unity Catalog metric views join each other — including cardinality and join condition — using `ALTER VIEW ... ADD CONSTRAINT ... RELATIONSHIP`, so AI assistants and BI tools automatically select the correct join path. Requires Databricks Runtime 19 or a SQL warehouse.

— Exact DDL syntax and runtime requirement given, directly runnable

Relationships between metric views (Bet…

02 **Native XML file format support (GA)** NEW 85

Databricks now supports ingesting, querying, and parsing XML files in batch or streaming workloads using Auto Loader, `read_files`, and `COPY INTO`, and parsing XML data in SQL via the `from_xml` and `schema_of_xml` functions. This is now generally available.

Extract structured fields from an XML column already stored in a Delta table using SQL.

```sql
SELECT from_xml(xml_col, schema_of_xml('<record><id>1</id>
<name>foo</name></record>')) AS parsed
FROM my_table;
```

— Names every ingestion path and function plus a runnable SQL example

Native XML file format support is now g…

03 **Real-time mode in Lakeflow pipelines (GA)** NEW 70

Lakeflow pipelines now support real-time mode on both classic compute and serverless compute, enabling operational workloads such as fraud detection with end-to-end latency as low as five milliseconds. This capability is now generally available.

— Describes scope and latency numbers but no config/flag or runnable step

Real-time mode in Lakeflow pipelines is…

THINNER COVERAGE BELOW

04 **Claude Code auto mode for more models** IMPROVED 40

Claude Code auto mode now works with both Anthropic and non-Anthropic models hosted by Databricks, reviewing tool calls to reduce routine permission prompts.

— Explains benefit but gives no config, command, or navigation path

Claude Code auto mode for Anthropic and…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Apps and Automation Service Integrations - October Release (2026-10-30) BLOG ↗

RANK

Sumo Logic is a cloud-native machine data analytics platform that collects, searches, and analyzes logs and metrics for security monitoring, troubleshooting, and operational insights.

> Sumo Logic added self-service feature opt-in/opt-out controls to the Organizations UI and expanded several Automation service integrations, including custom Jira priority support for Cloud SOAR playbooks, new ZeroFox CTI enrichment actions, and ThreatConnect tag modes, while deprecating legacy CrowdStrike Falcon Incident actions.

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

01 **Custom Jira priority support in Cloud SOAR playbooks** NEW

40

The Atlassian Jira Cloud integration and Atlassian Jira V2 integration both now add custom priority support for Cloud SOAR playbooks.

— Names two integrations but no config detail or UI path.

Apps and Automation Service Integration…

02 **Tag modes for ThreatConnect Update Indicators action** NEW

35

ThreatConnect V3 integration adds tag modes to the Update Indicators action.

— Names the specific action but not the tag mode options.

Apps and Automation Service Integration…

03 **CrowdStrike Falcon legacy Incident actions deprecated**

DEPRECATED

35

CrowdStrike Falcon legacy Incident actions are deprecated and may be removed in a future release.

— States deprecation but gives no migration path.

Apps and Automation Service Integration…

04 **Chats dashboard update in Claude Compliance** IMPROVED

25

Claude Compliance v2.0.6 updates the Chats dashboard with tool information.

— Version and dashboard named, but no detail on content.

Apps and Automation Service Integration…

05 **CTI enrichment actions in ZeroFox V2** NEW 20

ZeroFox V2 integration adds new CTI enrichment actions to the integration.

— Bare mention, no detail on what the actions do.

Apps and Automation Service Integration…

└──▷ **BREAKING ON UPGRADE**

! CrowdStrike Falcon legacy Incident actions are deprecated and may be removed in a future release.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Node v6.20.1: Leaked CA keys no longer trusted (2026-10-30) BLOG ↗

RANK

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

> Tenzir's v6.20.1 release trims small-file overhead in its Iceberg sink, pushes more filtering work to SentinelOne before data reaches the node, and stops trusting a set of leaked CA keys.

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

01 **Fewer small files from to_iceberg sink** IMPROVED 55

The `to_iceberg` operator now writes fewer small data files, reducing object-store overhead for Iceberg table sinks.

— Names the operator and benefit, no numbers or config shown

Tenzir Node v6.20.1: Leaked CA keys no …

02 **SentinelOne filter pushdown inside disjunctions** IMPROVED 40

Tenzir now pushes more SentinelOne filter conditions inside disjunctions to the SentinelOne side, reducing the amount of data transferred to the node.

— Explains mechanism but no query syntax or flag given

Tenzir Node v6.20.1: Leaked CA keys no …

03 **Leaked CA keys no longer trusted** BREAKING 10

This release stops trusting a set of leaked CA keys, per the release title; no further mechanism or affected-component detail was provided in the changelog.

— Only mentioned in title, no detail or scope given

Tenzir Node v6.20.1: Leaked CA keys no …

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

syzkaller

Feedback ⓘ changes since 2026-10-28 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds sandboxed build/git helpers, Markdown-linkified AI root cause analysis in the dashboard, and a linter rule for raw string literals.

syzkaller now exposes `LinkifyHTML` and `LinkifyMarkdown` functions in `pkg/html` to turn source file references (e.g. `fs/ext4/inode.c:2753`) into hyperlinks anchored to a given kernel repo and commit.

syzkaller's `syz-linter` now enforces raw string literals for all multi-line strings.

syzkaller adds `RunSandboxed` and `WriteSandboxFile` helpers in `pkg/osutil`, plus a `SandboxedGit` helper in `pkg/vcs`, so that git operations and kernel builds run under a sandbox.

syzkaller's `pkg/aflow` pipeline now runs git, kernel builds, and tools under a sandbox, and wipes cache temp dirs on startup.

syzkaller's dashboard now renders AI root cause analysis as Markdown.

NEW DETECTION VMX MSR load/store list syscall descriptions for x86 SYZOS

UPDATED DETECTION AF_VSOCK syscall descriptions for Linux

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

oxo

Feedback ⓘ v2.22.0 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> oxo adds support for scanning ASN (Autonomous System Number) targets, expanding the range of scan target types it can receive and run from the scanning engine.

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

01 **ASN scan target support** NEW 40

oxo now runs ASN (Autonomous System Number) scans received from the scanning engine, adding ASN to the list of supported scan target types.

— Names new target type but no command or config detail v2.22.0

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

ITarian

Feedback ⓘ

From Powerful to Unstoppable: The ITarian 2026–2027 Roadmap (2026-10-30) BLOG ↗

RANK

Streamline IT operations with ITarian's all-in-one IT management software. Automate tasks, enhance security, and boost efficiency.

> ITarian's 2026–2027 roadmap adds exposure management, ransomware protection, and security-awareness training alongside broader cross-platform endpoint, scripting, and MDR integration coverage.

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

01 **Cross-platform scripting and monitoring for Linux and Mac**

NEW 45

ITarian now supports `Shell/Bash` and `Python 3` scripting on Linux endpoints alongside existing Windows scripting, and adds Linux and Mac monitors for cross-platform endpoint visibility.

– Names scripting languages and platforms but no workflow detail

From Powerful to Unstoppable: The ITari…

02 **MDR integrations for Entra, Google Workspace, and CrowdStrike**

NEW 40

ITarian adds MDR integrations with Microsoft Entra, Google Workspace, and CrowdStrike.

– Names three integration targets but no setup steps

From Powerful to Unstoppable: The ITari…

03 **Automatic EXDR install and auto-update** IMPROVED 40

ITarian adds automatic EXDR install and auto-update for endpoint detection and response deployment.

– Describes the automation but not rollout mechanics

From Powerful to Unstoppable: The ITari…

04 **Default-on network share ransomware protection** NEW 35

ITarian adds network share ransomware protection, enabled by default.

– States default-on behavior but no config or scope detail

From Powerful to Unstoppable: The ITari…

05 **CTEM exposure management module** NEW 30

ITarian adds a CTEM (Continuous Threat Exposure Management) module for proactive exposure management.

— Names the module but no mechanism or workflow detail

From Powerful to Unstoppable: The ITari…

06 **SOC 2 and HIPAA compliance gap assessments** NEW 30

ITarian adds SOC 2 and HIPAA gap assessments to its compliance tooling.

— Names two frameworks but no assessment process detail

From Powerful to Unstoppable: The ITari…

07 **IP-based login restrictions** NEW 30

ITarian adds IP-based login restrictions for access control.

— Clear access-control concept but no config path given

From Powerful to Unstoppable: The ITari…

08 **Windows ARM64 device support** NEW 30

ITarian adds Windows ARM64 device support.

— Bare platform-support announcement, no further detail

From Powerful to Unstoppable: The ITari…

09 **Phishing simulation and cyber awareness training** NEW 25

ITarian adds Phishing Simulation and Cyber Awareness Training capabilities.

— Names the capability with no further mechanism given

From Powerful to Unstoppable: The ITari…

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ What's New in Iru | October 30th NOTES ↗ RANK

Iru is a device-centric platform that unifies security, IT, and compliance around devices and everything they touch.

> Iru extended Privacy Library coverage to Mac, added in-app Android managed Google domain migration, and improved the SSO sign-in experience.

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

01 **In-app Android managed Google domain migration** NEW 40

Admins can upgrade Android devices to a managed Google domain directly within Iru.

— Clear action but no steps or limits given What's New in Iru | October 30th

02 **Privacy Library Item coverage for Mac** NEW 30

Privacy Library Item coverage now extends to Mac computers.



— Names the platform but no mechanism or scope detail What's New in Iru | October 30th

03 **SSO sign-in experience upgrades** IMPROVED 20

Single Sign-On received upgrades intended to improve the user sign-in experience.

— Vague improvement with no specifics on what changed What's New in Iru | October 30th

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Backup & Recovery

Acronis Cyber Infrastructure

Feedback ⓘ

changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Acronis Cyber Infrastructure is hyperconverged software that pools storage, compute, and networking on commodity servers, often as a backup storage target.

> Acronis Cyber Infrastructure adds a command-line tool for importing StoreOnce datastores into the backup gateway, enriches workload reports with application version data, and speeds up S3 bucket cleanup with a delete-by-prefix option.

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

01 **Delete-by-prefix for S3 bucket cleanup** NEW 55

S3 bucket deletion now supports a delete-by-prefix operation, intended for buckets holding a very large number of objects.

– Names the operation and use case but no exact API/command changelog-20261031-5816b1ba

02 **CLI tool for importing StoreOnce datastores** NEW 45

A new command-line tool lets administrators import an existing StoreOnce datastore into the backup gateway.

– Names the capability but not the exact command changelog-20261031-5816b1ba

03 **App name and version in workload reports** IMPROVED 45

Workload reports now include the product name and version of applications discovered on Windows VMs.

– Describes the change but no report location or format detail

changelog-20261031-5816b1ba

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Codezero

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Codezero is a aI agents, tools, apps, and pipelines use dozens of credentials with every interaction. They don't need to hold a single one.

> Codezero's Cordon shipped a CA bundle refresh command to recover from the eMudhra certificate authority compromise, hardened its TLS verifier against the compromised intermediates, broadened dependency audits, and tightened Node.js and proxy support requirements.

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

01 **CA bundle refresh command** NEW 90

New `cordon env --refresh-ca-bundle` command rebuilds a configuration's combined CA bundle from its installed public-CA source and existing local Cordon CA, preserving the CA certificate, private key, and configuration. Intended for use after the eMudhra CA compromise to re-establish trust without rotating the local CA or rerunning `cordon trust`.

After the eMudhra CA compromise, refresh each scope's CA bundle to re-establish trust without rotating the local CA or rerunning `cordon trust`.

```
$ cordon env --refresh-ca-bundle > /dev/null
```

— Named flag, exact behaviour, and runnable example given. changelog-20261031-fddd5f41

02 **Dropped Node.js 20 support** BREAKING 60

The npm package now requires Node.js 22.23.4+, 24.22.1+, or 26.12.1+ within those major versions; Node.js 20 is no longer supported.

— Exact version thresholds given, clear migration implication.

changelog-20261031-fddd5f41

THINNER COVERAGE BELOW

03 **Axios helper drops PAC and SOCKS proxy support** BREAKING 53

The Axios helper no longer accepts PAC or SOCKS proxy URLs; only HTTP and HTTPS proxy URLs are supported.

— Specific before/after but no reconfiguration steps shown. changelog-20261031-fddd5f41

04 **TLS verifier distrusts eMudhra intermediates** IMPROVED

52

The TLS verifier now explicitly distrusts the compromised eMudhra intermediate public keys, including cross-signed chains and installations with older system roots.

— Mechanism and scope named but no user-facing action. changelog-20261031-fddd5f41

05 **Broader dependency auditing in CI** IMPROVED

40

Dependency audits now include development and optional dependencies and run daily on main, with release builds preserving locked dependency resolution.

— Describes scope and cadence but no reader action. changelog-20261031-fddd5f41

└──▷ **BREAKING ON UPGRADE**

! The npm package now requires Node.js 22.23.4+, 24.22.1+, or 26.12.1+ within those major versions; Node.js 20 is no longer supported.

! The Axios helper no longer accepts PAC or SOCKS proxy URLs; only HTTP and HTTPS proxy URLs are supported.

WAS THIS USEFUL?  

Headscale

Feedback ⓘ v0.29.7 NOTES ↗ RANK

An open source, self-hosted implementation of the Tailscale control server

> Headscale v0.29.7 adds a way to lock down Tailscale SSH agent and port forwarding, closes a metrics exposure gap, and hardens registration, key-expiry, and client-report handling against abuse and resource exhaustion.

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

01 Disable SSH agent and port forwarding per node NEW

85

Setting `TS_SSH_DISABLE_FORWARDING=1` in the `tailscaled` service environment disables SSH agent and port forwarding for all Tailscale SSH sessions on that node, and works even on older Tailscale clients.

Lock down SSH agent and port forwarding on a node exposed to untrusted users via Tailscale SSH.

text

```
# In the tailscaled systemd service override or environment file:
TS_SSH_DISABLE_FORWARDING=1
```

— Exact env var and config file given, mechanism and scope explained. v0.29.7

02 Metrics no longer reachable by unregistered clients IMPROVED

85

Metrics are no longer exposed to unregistered clients through the Noise protocol and are only available on the configured `metrics_listen_addr`; setting `metrics_listen_addr` to empty disables the metrics listener entirely.

Restrict Headscale metrics so they are only reachable on the local management interface, not through the Noise protocol.

yaml

```
metrics_listen_addr: '127.0.0.1:9090'
```

— Names the config key and gives a runnable config example. v0.29.7

03 Key expiry value validation NEW

75

Key expiry values outside the years 0–9999 UTC are now refused at registration, on the OIDC token-expiry path, and in `headscale nodes expire` / `POST /api/v1/node/{node_id}/expire`, which now return HTTP 400 for invalid values.

— Names CLI command and exact API endpoint with status code. v0.29.7

04 **Client report size limits** NEW 70

Client reports are now bounded to the first 256 endpoints and 4096 announced routes, with oversized hostnames, requested tags, or other metadata returning a registration error or HTTP 413 on map updates, and zoned IPv6 endpoints refused; at startup Headscale trims stored hostnames, endpoints, and routes to these limits and removes zoned endpoints.

— Concrete numeric limits and error codes, but no reader action. v0.29.7

05 **Bounded capacity for pending registrations** NEW 65

Pending interactive registrations and OIDC logins are kept at capacity rather than evicted; new requests beyond capacity get a retry response, registered nodes keep capacity reserved for SSH checks, and stored pending registration data (tags and expiry retained) is capped at `64 KiB`, with oversized interactive registrations rejected.

— Mechanism and exact 64 KiB cap given, but no direct command. v0.29.7

06 **SSH policy acceptEnv fix** IMPROVED 60

SSH policy `acceptEnv` now works correctly, allowing configured environment variables to reach SSH sessions; upgrade SSH destination nodes to Tailscale v1.104.4 or newer before using `acceptEnv`, since older clients can disclose forwarded values.

— Names the policy field and gives a concrete upgrade requirement. v0.29.7

THINNER COVERAGE BELOW

07 **Tighter re-registration rules for tagged nodes** BREAKING 55

Re-registering an expired tagged node now requires a valid pre-auth key, and using another user's key or login on a tagged machine now creates a new device for that user instead of re-authenticating the tagged node.

— Behavior change described but no command or endpoint given. v0.29.7

08 **Idle connection timeout for Noise tunnels** NEW 30

A stalled request body, or a Noise tunnel with no open request, is now closed after 30 seconds.

— Brief behavioral note with a single number, no further detail. v0.29.7

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

RoboForm

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

RoboForm is a password manager and form-filling tool that securely stores login credentials and automatically fills web forms across browsers and devices.

> RoboForm shipped a handful of small refinements this window, including automatic clipboard clearing in the Editor, broader Start Page search coverage, and updates to its encryption and desktop unlock authentication flows.

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

01 **Clipboard auto-clear after Editor copy** IMPROVED 55

RoboForm's Editor now auto-clears the clipboard two minutes after copying data with the Editor's Copy button, reducing the risk of sensitive data lingering in the clipboard.

– Names the exact trigger and timing, but no config option mentioned.

changelog-20261031-25cd6930

02 **Folders included in Start Page search** NEW 35

RoboForm now includes folders in Start Page search results, in addition to whatever was previously searchable.

– States what changed but no detail on search mechanics. changelog-20261031-25cd6930

03 **New authentication protocol for Desktop Unlock UI** IMPROVED 35

RoboForm updates the Desktop Unlock UI to use a new authentication protocol, though the protocol itself is not identified.

– Identifies the UI surface changed but not the protocol used.

changelog-20261031-25cd6930

04 **Updated encryption and authentication settings** IMPROVED 20

RoboForm adds support for updated encryption and authentication settings, though the specific algorithms or protocols are not named.

– Vague wording with no named settings or protocols. changelog-20261031-25cd6930

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2026 (2026-10-30) BLOG ↗

RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud's biggest additions this release center on Workflows: natural-language control via MCP-enabled AI clients, OAuth 2.0-secured connectors, and new automation templates, alongside App Catalog update-notification controls, device Agent Uninstall Protection with maintenance tokens, and brief mentions of SIEM/OCSF export, AI agent management APIs, and custom banned passwords.

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

01 **App Catalog update notification controls** NEW 78

Admins can control App Catalog update notifications by selecting Disable end user update notifications when adding or editing an application, and can allow users to start Required app updates from the Self Service App Catalog.

— Exact toggle name and self-service path given Release Notes 2026 (2026-10-30)

02 **OAuth 2.0 authentication for Workflow connectors** NEW 70

Workflow connectors now support OAuth 2.0 authentication, including Client Credentials and Authorization Code with PKCE, with credentials stored in Secrets Manager and access tokens refreshed automatically.

— Named auth flows and storage, but no config key or endpoint shown

Release Notes 2026 (2026-10-30)

03 **Workflow management via MCP-enabled AI clients** NEW 63

Admins can create, modify, duplicate, and troubleshoot JumpCloud Workflows from any MCP-enabled AI client using natural language, including triggering workflows and reviewing per-node outcomes, through the existing JumpCloud MCP Server connection.

— Mechanism named (MCP Server connection) but no exact command shown

Release Notes 2026 (2026-10-30)

04 **New Workflow templates for common automation scenarios** NEW 60

New Workflow templates cover automation scenarios including auto-suspending inactive users, blocking newly discovered generative AI apps, revoking Mac admin privileges on enrollment, detecting sensitive software installs, and managing pre-hire user groups.

.
.
.
.

— Five scenarios named, no exact setup steps given Release Notes 2026 (2026-10-30)

THINNER COVERAGE BELOW

05 **SIEM integration with OCSF export** NEW 25

JumpCloud adds SIEM integration with OCSF export, per the release summary, though no further mechanism or configuration detail is provided.

— Named only in summary, no detail or steps given Release Notes 2026 (2026-10-30)

06 **AI agent management APIs** NEW 10

JumpCloud adds AI agent management APIs, as noted in the release summary, with no further detail on scope or endpoints provided.

— Named only in summary, no detail given Release Notes 2026 (2026-10-30)

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Multi-Factor Authentication: Exponential Backoff for Phone MFA SMS Delivery (2026-10-30)

BLOG ↗ RANK

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

> Auth0 introduced SMS Fraud Protection for Phone MFA on Enterprise plans, replacing static rate limits with exponential backoff and adding audit logging for throttled attempts.

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

01 **Exponential backoff for Phone MFA SMS** NEW 70

Auth0 adds SMS Fraud Protection with exponential backoff for Phone MFA on Enterprise plans, replacing static rate limits with progressive delays. An initial buffer of immediate retry attempts accommodates typos or delivery delays before progressive wait times take effect, and throttled SMS send attempts are logged as `limit_phone` tenant log events, including remaining cooldown durations, for audit visibility into fraud-throttling activity.

— Names log event and mechanism but no config/flag to enable it

Multi-Factor Authentication: Exponentia…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ

Read Application Logs with the Backend API (2026-10-30) BLOG ↗ RANK

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

> Clerk opened up application log access via a public Backend API endpoint, enabling programmatic querying of log events for custom tooling or AI agents.

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

01 **Application logs exposed via Backend API** NEW 75

The `GET /v1/logs` Backend API endpoint is now public, allowing application log events to be queried programmatically from custom code or an AI agent, e.g. `curl -X GET 'https://api.clerk.com/v1/logs' -H 'Authorization: Bearer <your_secret_key>'`.

Fetch recent application log events from Clerk to feed into an AI agent or custom monitoring pipeline.

```
$ curl -X GET 'https://api.clerk.com/v1/logs' \
  -H 'Authorization: Bearer <your_secret_key>'
```

— Names exact endpoint and runnable curl example, but no response schema or filters detailed.

Read Application Logs with the Backend …

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Network, Ory Hydra, Ory Kratos, Ory Keto, Ory Oathkeeper, Ory Polis, Ory Elements v2

NOTES ↗ RANK

Discover the latest features, improvements, and updates to Ory.

> Ory shipped per-stream event type filtering, letting event streams be scoped to specific event types rather than delivering everything.

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

01 **Event type filtering for event streams** NEW 60

Event streams can now be filtered to a specific set of event types by setting `event_types` on the stream, instead of receiving all events.

— Names exact config field but no further mechanism detail.

Ory Network, Ory Hydra, Ory Kratos, Ory…

WAS THIS USEFUL? 👍 👎

Omada Identity

Feedback ⓘ 29 October 2026 (2026-10-29) BLOG ↗

RANK

Omada Identity is an identity governance and administration platform that manages user access, entitlements, compliance certifications, and access requests across enterprise systems.

> Omada Identity's sole update this window improves operational visibility into the Cloud Application Gateway worker heartbeats shown in the Operations dashboard.

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

01 **Instance IDs in System Heartbeats widget** IMPROVED 55

The System Heartbeats widget in the Operations dashboard now displays the instance ID next to each worker service, allowing administrators to distinguish between customer-hosted and Omada-hosted Cloud Application Gateway (CAG) instances when both send heartbeats.

— Names exact widget and dashboard but no further mechanism. 29 October 2026 (2026-10-29)

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Varonis

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Varonis is the leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis shipped a new Entity ID Lookup API, CSV export support, expanded EDM classification with offline hashing, and several workflow conveniences including Microsoft Main link monitoring, in-UI classification pattern management, data source collections, and automation cloning.

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

01 **Entity ID Lookup API** NEW 75

The new Entity ID Lookup API translates physical folder paths into entity IDs required by the Search and Export APIs, supporting exact path resolution, bulk lookups, path discovery, and subtree resolution.

– Describes four lookup modes but no endpoint path or method given

changelog-20261031-65059257

02 **EDM classification for tabular files with offline hashing** NEW 70

Exact Data Match (EDM) classification now applies to tabular files using the same policies as non-tabular data, and EDM source data can be prepared offline by hashing files on disconnected systems and uploading only the hashed package from a connected environment.

– Workflow mechanism explained, no exact command or limit given

changelog-20261031-65059257

THINNER COVERAGE BELOW

03 **CSV export format for Export API** NEW 50

The Export API now supports exporting data in `CSV` format in addition to `AVRO`, simplifying analysis in spreadsheet and reporting tools without requiring file conversion.

– Names the formats but gives no endpoint or parameter changelog-20261031-65059257

04 **Classification pattern management in Web Interface** NEW 50

Classification patterns can now be created, edited, tested, and managed directly in the Varonis Web Interface, integrated with classification policies and settings.

– Clear UI location but no detail on testing mechanism changelog-20261031-65059257

05 **Microsoft Main link monitoring** NEW 45

Varonis now monitors and manages Microsoft Main links within SharePoint Online and OneDrive, extending analytics and policy workflow coverage to supported Main link types.

— States scope but no specifics on policy actions changelog-20261031-65059257

06 **Named data source collections** NEW 40

Multiple data sources can now be organized into named, reusable collections that can be applied across analytics, dashboards, and policies.

— Describes capability generically without limits or UI path changelog-20261031-65059257

07 **Clone data protection automations** NEW 40

Existing data protection automations can be cloned to reuse configurations instead of rebuilding them from scratch.

— Simple capability stated without mechanism or scope changelog-20261031-65059257

WAS THIS USEFUL?  

Zama fhEVM

Feedback ⓘ v0.14.5 NOTES ↗

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

fhEVM adds a relayer query option to the SDK and removes transient allowance for inputs to sender in the Solidity library.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.14.5 https://github.com/zama-ai/fhevm.git
# already have the repo? check out this version:
$ git checkout v0.14.5
```

Add a relayer query option to the SDK via the backported relayer query option (#4544), giving developers control over relayer selection when submitting encrypted transactions.

NEW SUPPORTED PLATFORM Remove transient allowance for inputs to sender in the Solidity library, changing how input permissions are granted on-chain.

└──▷ BREAKING ON UPGRADE

! Transient allowance for inputs to sender has been removed from the Solidity library (`library-solidity`); contracts relying on this allowance behaviour will no longer work as before.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 7.0.0 (2026-10-31) BLOG ↗ RANK

RegScale is a compliance automation platform that helps organizations manage security controls, risk, and regulatory frameworks like NIST, FedRAMP, and SOC 2 through continuous monitoring.

> RegScale CLI 7.0.0 adds optional OpenTelemetry and Datadog observability across all CLI commands.

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

01 **OpenTelemetry/Datadog observability for the CLI** NEW

70

All CLI commands, outbound HTTP/GraphQL requests, sync item counters, and secret-redacted structured logs can now emit telemetry over OTLP/gRPC. Enable it by installing the `telemetry` extra and configuring via the `REGSCALE_OTEL` environment variable.

— Names env var and extra but no further config detail CLI 7.0.0 (2026-10-31)

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

Qualys Policy Compliance

Feedback ⓘ

changelog · Oct 30, 2026 BLOG ↗

Qualys Policy Compliance is a cloud-based security module that audits IT systems against configuration benchmarks and regulatory standards, identifying misconfigurations and reporting compliance gaps.

NEW COMPLIANCE FRAMEWORK CISA AA26-194A out-of-the-box policy for Cisco router configuration hardening

WAS THIS USEFUL? 👍 👎

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Oct 29th, 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.

UPDATED INTEGRATION OpenAI API Platform integration expanded data collection

UPDATED INTEGRATION Google Cloud Platform integration expanded data collection

UPDATED INTEGRATION BitSight integration expanded data collection

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 2.2.0 (2026-10-30) BLOG ↗ LATEST STABLE

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

Runlayer adds `update_skill_shares` and `list_skill_shares` MCP tools for managing skill access, and metric definitions on Control and Usage dashboards.

Share skills with teams or people over the Runlayer MCP using the new `update_skill_shares` tool, which sets permissions or removes access under the same rules as the skill's Share dialog.

Inspect current skill access over the Runlayer MCP using the new `list_skill_shares` tool, with share counts also surfaced in `list_skills`.

Every number card on Control and Usage now shows an ⓘ metric definition explaining what it counts and what the trend arrow compares to.

NEW INTEGRATION AWS MCP Server connector

UPDATED INTEGRATION Shadow AI inventory — VS Code MCP server discovery

UPDATED INTEGRATION Shadow AI inventory — GitHub Copilot CLI MCP server discovery

UPDATED INTEGRATION AI Watch Enforce — GitHub Copilot CLI MCP server policy enforcement

UPDATED INTEGRATION AI Watch Enforce — VS Code legacy Local chat harness blocking

▾ 1 more release in this issue · 2026-10-30

> 2.1.0 (2026-10-30) BLOG ↗ STABLE
>
> Runlayer adds per-run LLM spend limits ($250 default), agent-change approval flows in chat, Kubernetes audit consumer shard leasing, and in-conversation skill editing.
>
> Agents now stop a run once it has spent $250 (or the remaining daily/weekly budget, whichever is lower) when LLM traffic flows through the Runlayer LLM gateway, with the run failing with a named limit error and the agent owner receiving a Slack DM when Budget alerts are on.
>
> Runlayer now asks the person in a conversation to approve any agent-initiated change to its own setup (connectors, skills, instructions, subagents, schedules, triggers, webhooks, Slack settings, channel monitors, or adding a workspace connector) before applying it, including changes requested over iMessage and approved in Runlayer.

The Kubernetes audit consumer now leases its Kinesis shards using DynamoDB checkpointing; custom `{instance}-audit-consumer` IAM policies must allow `dynamodb:Query` and `dynamodb:UpdateItem` on the checkpoint table, and `dynamodb:DeleteItem` scoped to `dynamodb:LeadingKeys = workers#*`.

Handing a task to another agent now returns in about one second instead of about 20 seconds.

Agents can now create and update skills directly in conversation without requiring an explicit Runlayer connector or a ZIP file upload.

└──▷ **BREAKING ON UPGRADE**

- ! If you wrote a custom `{instance}-audit-consumer` IAM policy instead of using the published sample, you must now add `dynamodb:Query` and `dynamodb:UpdateItem` on the checkpoint table and `dynamodb:DeleteItem` scoped to `dynamodb:LeadingKeys = workers#*`, or the consumer will log `kinesis_consumer_missing_permissions` at startup and exit.
- ! Legacy macOS MDM scripts (AI Watch and Hooks) must be regenerated and redeployed: they now run `uv` and the Runlayer CLI as the signed-in user instead of root and keep the enrollment key off the command line.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ AI / LLM Gateways

Portkey

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Democratize and productionize Gen AI across your entire org with Portkey

> Portkey expanded its guardrail system to cover MCP tool calls and introduced a new pre-model lifecycle stage for inspecting traffic before it reaches a model.

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

01 **Guardrails extended to MCP tool calls** NEW 35

Portkey now applies guardrails to MCP tool calls, extending coverage beyond standard model request/response traffic.

— Names the surface (MCP tool calls) but no mechanism detail. changelog-20261031-d26aea6d

02 **Pre-model traffic inspection lifecycle stage** NEW 23

Portkey adds a new lifecycle stage to inspect traffic before it reaches a model.

— Bare description with no mechanism or configuration detail. changelog-20261031-d26aea6d

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Runtime Defense & Guardrails

Lakera

Feedback ⓘ changelog · Oct 31, 2026 CHANGELOG ↗ RANK

Lakera is an AI security platform that protects large language model applications from prompt injection, data leakage, and harmful content through real-time guardrails and APIs.

> Lakera added a request latency percentiles chart to its Analytics page for monitoring Guard API performance over time.

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

01 **Request latency percentiles chart in Analytics** NEW

62

The Analytics page now plots `p25`, `p50` (median), `p75`, and `p95` latency of Guard API requests over time, with hover-to-inspect values per time bucket and click-to-toggle visibility per percentile via the legend.

Spot tail-latency spikes in Guard API traffic by checking p95 on the Analytics page.

📍 In the Lakera console, go to Analytics › Request Latency Percentiles. Hover over any point to read p25/p50/p75/p95 values for that time bucket. Click a legend entry (e.g. `95th`) to isolate that percentile on the chart.

— Describes exact metrics and UI interactions but only one chart.

changelog-20261031-3b91e91c

WAS THIS USEFUL? 👍 👎

◆ AI Red Teaming & Testing

humanbound

Feedback ⓘ v2.14.0 NOTES ↗ RANK

Open-source adversarial testing engine, SDK, and CLI for AI agents. Runs locally or against the Humanbound Platform.

> humanbound's v2.14.0 release hardens arena gateway token handling with zero-downtime rotation and default masking, and adds structured scan-outcome reporting for `hb connect`.

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

01 **Zero-downtime arena token rotation** NEW 87

`hb arena token --rotate` stops the running arena gateway, swaps in a new access token, and restarts it; only the new token is written to stdout, so the old one is immediately invalidated.

Rotate the arena gateway token in-place and immediately inject it into subsequent commands without downtime.

```
$ TOKEN=$(hb arena token --rotate)
hb test --target arena://my-agent
```

— Exact command with full before/after mechanism and example. v2.14.0

02 **Default token masking and reveal in arena endpoint** BREAKING 83

`hb arena endpoint` now masks the access token by default so its output is safe to paste or show on screen; `hb arena endpoint --show-token` reveals the real unmasked token when needed. This is a breaking change — `hb arena endpoint` no longer prints the token, so scripts that read it from that command's output must switch to `hb arena token`.

Inspect the real unmasked gateway token when you need to paste it into an external config.

```
$ hb arena endpoint --show-token
```

— Names exact flags/commands plus the required migration step. v2.14.0

03 **Scan outcome reporting in hb connect** NEW 71

`hb connect` now reports agent scan outcomes by sending `connect_scan_started` before `POST /scan`, then `scan_completed` or `scan_failed`, carrying connection type and duration, and on failure an error type of `timeout`, `http_error`, or `network`.

— Names event names, endpoint, and error types but no user-facing command. v2.14.0

THINNER COVERAGE BELOW

04 **Non-loopback gateway exposure warning** IMPROVED 42

`hb arena serve --host <non-loopback>` now warns that the access token is the only network protection in place and suggests rotating it afterwards.

— Names the flag but describes only a warning message, no further mechanism. v2.14.0

└──▷ **BREAKING ON UPGRADE**

! `hb arena endpoint` no longer prints the access token — scripts that read the token from its output must switch to `hb arena token`.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

Workers, BYOIP, Bots - Workers subrequests now preserve context for BYOIP destinations (

BLOG ↗ RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

> Cloudflare's biggest change this window is a breaking shift in how Workers subrequests to BYOIP destinations report client IPs and bot signals, alongside new visibility tools for SCIM provisioning and account audit activity, a configurable delay for Durable Object code updates, and a hard size limit coming to D1 query results.

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

01 **BYOIP Workers subrequest context preservation** BREAKING

85

Workers subrequests to BYOIP-proxied hostnames now preserve the `CF-Worker` header and `cf.worker.upstream_zone` field, enabling consistent security rule evaluation across same-zone and cross-zone subrequests. Bot Management fields on BYOIP destination zones (bot scores, verified bot status, JA3/JA4 fingerprints) now reflect the original end-user connection rather than the internal Cloudflare service connection. Cross-zone Workers subrequests to BYOIP destinations now set `CF-Connecting-IP` to the Workers client IP `2a06:98c0:3600::103` instead of the original client IP, while same-zone subrequests continue to reflect the original client IP; any security rule or application logic reading `CF-Connecting-IP` for cross-zone subrequest traffic must be updated.

— Names exact headers, field, and IP with migration note

Workers, BYOIP, Bots - Workers subreque…

02 **Durable Objects deferred code updates** IMPROVED

65

Active Durable Objects now defer code updates by default, allowing up to 5 minutes to finish work before adopting new code; the delay is configurable or can be bypassed for immediate updates.

— States default behaviour and timing but not the config key product docs

03 **D1 50 MB query result limit** BREAKING

60

Beginning November 4, 2026, D1 queries that return a result larger than 50 MB will fail and return an error.

— Exact limit and effective date, no migration guidance product docs

THINNER COVERAGE BELOW

04 **Keyless SSL TLS 1.3 negotiation** IMPROVED 55

Keyless SSL now negotiates TLS 1.3 for connections that use ECDSA certificates when TLS 1.3 is enabled on the zone.

— Specific condition given, no setup steps product docs

05 **SCIM population viewer in Cloudflare One** NEW 45

Cloudflare One adds a SCIM population viewer that lets administrators search SCIM users and groups, inspect resource attributes, and investigate provisioning issues.

— Names the feature and its purpose, no navigation path product docs

06 **Audit Logs v2 account analytics** NEW 45

Audit Logs v2 adds account analytics to analyze account activity over time and investigate actions, actors, products, and failures.

.

— Describes capability with screenshot, no exact UI steps product docs

07 **Regional Gateway HTTP inspection for WAN traffic** NEW 40

Cloudflare WAN traffic can now be routed through a selected Gateway processing region for TLS decryption and HTTP inspection.

— Names mechanism but not configuration steps product docs

08 **Classic DNS records page retired** BREAKING 30

The classic DNS records page has been retired; all users now use the new DNS records experience.

— Minimal detail on what changed in the new experience product docs

09 **Browser Run default browser upgrade** IMPROVED 30

Browser Run now uses Chrome 152 as its default browser.

— Bare version bump with no further detail product docs

└──▷ BREAKING ON UPGRADE

! Beginning November 4, 2026, D1 queries that return a result larger than 50 MB will fail and return an error.

! The classic DNS records page has been retired; all users now use the new DNS records experience.

! Cross-zone Workers subrequests to Cloudflare-proxied BYOIP hostnames now set `CF-Connecting-IP` to `2a06:98c0:3600::103` instead of the original client IP. Any security rule or application logic on the destination zone that reads `CF-Connecting-IP` to allow/block specific IPs for cross-zone subrequest traffic must be updated.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 31, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue