

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 18, 2026 · EVERY DAY

EDITIONS **tail** | grep | head | ls | diff | uniq

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 24 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. Qualys TruRisk Eliminate 2.0 moves patching off the ticket queue into continuous ring-based rollout with
 vim ? help her coverage for agentless edge devices. python-pipx adds a

release-age `--cooldown` so freshly published — and possibly compromised — package versions can't install. Finite State can now diff two Project Versions by findings and components over the API.

PROTECT

Patch continuously in rings instead of cutting a ticket per CVE

TruRisk Eliminate 2.0 pairs the TruConfidence orchestration engine with the autonomous Agent Sara to drive remediation as a continuous, ring-based rollout rather than a ticket queued to an owner who may never act on it. Scanner-based coverage picks up edge devices that never got an agent, which is where unpatched exposure usually survives longest.

[Qualys Patch Management](#)

PROTECT

Refuse to install a package version that is only hours old

A `--cooldown` / `PIPX_COOLDOWN` policy holds back releases until they have aged, which is the practical defence against a hijacked maintainer account pushing a malicious version that gets yanked within a day. `--index-url` / `--pip-args` / `PIP_*` support plus `pipx manifest` with [PEP 751](#) lock files means installs can be pinned, reproducible and sourced from an internal index rather than resolved fresh off PyPI.

[python-pipx](#)

IDENTIFY

Diff two builds of a product and get just the findings that changed

A public API endpoint compares two Project Versions by findings and components, so a release gate can ask what this firmware build introduced rather than re-reviewing the whole [SBOM](#) and findings list. That turns a manual spreadsheet reconciliation between versions into a scriptable check.

[Finite State Platform](#)

GOVERN

Stop an agent mid-reply and keep red-team actions inside approved scope

Runlayer adds a Slack stop button that halts an in-progress agent reply, plus Google [ADK](#) session telemetry for Python agents — a kill switch and a record of what the agent did, which is what oversight of autonomous actions actually requires. Decepticon v1.5.0 gates Red actions behind a

stricter scope approval and adds a Blue Cell monitoring layer, so an offensive agent cannot wander off the authorised target list unsupervised.

[Runlayer](#) · [Decepticon](#)

OFFENSIVE SECURITY

PurpleAILAB Decepticon

[Feedback](#)

v1.5.0

[NOTES ↗](#)

RANK

Autonomous Hacking Agent for Red Team

Decepticon's v1.5.0 release adds session resume with cross-session queuing, a Blue Cell monitoring layer, a stricter scope-approval gate for Red actions, and optional vendor-neutral OpenAI-compatible gateway support, alongside several smaller diagnostics and UX additions.

↳ **WHAT SHIPPED** • **10 FEATURES** most completely described first [? what's the number?](#)

01 Session resume with cross-session instruction queue

NEW

80

Adds `--resume` / `--continue` flags to restore scoped conversations and pending questions across sessions. Failed runs retain pending input and queues are isolated per session, and a new `/queue` command manages waiting instructions between sessions.

Resume a prior engagement mid-session without losing pending questions or queued instructions.

```
$ decepticon --resume
```

— Names exact flags and command with a runnable example

v1.5.0

02 Scope import and explicit Red approval gate

BREAKING

70

Adds `/scope import <file> --confirm` to import reviewed scope, and now requires explicit `/mode red approve` before granting Red execution authority. Unscoped legacy conversations are now reachable only under `/sessions legacy` and require explicit linking to their engagement.

Import a pre-reviewed scope file and explicitly approve Red execution authority before any offensive action runs.

```
$ /scope import targets.txt --confirm
/mode red approve
```

— Named commands and behavior change, includes example

v1.5.0

03 Read-only diagnostics via doctor command NEW

60

Adds `decepticon doctor --json` for read-only diagnostics output.

Pull structured diagnostics from a running Decepticon instance for automated health checks or CI pipelines.

```
$ decepticon doctor --json
```

– Exact command with example, but thin on mechanism v1.5.0

THINNER COVERAGE BELOW

04 Blue Cell monitoring commands and state NEW

55

Adds `/alerts` and `/status` commands for Blue Cell investigation details and connectivity monitoring during Red conversations, plus an independent Blue monitoring state that surfaces compact notifications during Red conversations.

– Named commands but no runnable example given v1.5.0

05 Server, launcher and Docker image sync required

BREAKING

40

The new validation and role-purpose endpoints require the updated server — launcher and Docker images/configuration must be updated together.

– Mentions endpoints generically without exact paths v1.5.0

06 Conversation export to Markdown or JSON NEW

35

Adds conversation export to Markdown or JSON formats.

– Names formats but no command or flag given v1.5.0

07 Vendor-neutral OpenAI-compatible gateway support NEW

35

Adds optional vendor-neutral OpenAI-compatible gateway support.

– No configuration detail or endpoint named v1.5.0

08 Name-only engagement creation opens Red directly IMPROVED

35

Name-only engagement creation now opens Red directly without requiring a plan.

– Describes a workflow change but no command shown v1.5.0

09 Workflow checkpoints retain pinned source

BREAKING

Existing workflow checkpoints retain their pinned source; to change a workflow, a new execution conversation must be created.

– States behavior change but lacks concrete migration steps v1.5.0

10 Persistent dashboard light/dark theme NEW

15

Adds persistent light/dark dashboard preference.

– Bare feature name with no further detail v1.5.0

↳ **BREAKING ON UPGRADE**

- ! The new validation and role-purpose endpoints require the updated server — launcher and Docker images/configuration must be updated together.
- ! Unscoped legacy conversations are now available only under `/sessions legacy` and require explicit linking to their engagement.
- ! Existing workflow checkpoints retain their pinned source; to change a workflow, a new execution conversation must be created.

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

Finite State Platform

[Feedback](#)[changelog · Oct 18, 2026](#)[CHANGELOG](#)[RANK](#)

Finite State Platform is a firmware and software supply chain security tool that analyzes binaries and embedded devices to identify vulnerabilities, SBOM components, and risks.

Finite State Platform shipped a new public API endpoint for comparing two Project Versions by findings and components, alongside performance and usability refinements to the Projects API and Policy Overview UI.

←→ **WHAT SHIPPED · 3 FEATURES** most completely described first [? what's the number?](#)

01 Compare Project Versions by findings and components

NEW

90

New **GET**`/api/public/v0/versions/{baseVersionId}/compare/{targetVersionId}`

endpoint compares two Project Versions, surfacing which Findings and Components changed between them. Comparisons can return a 503 busy status, and callers should back off and retry (e.g. respecting a short delay) until the comparison completes.

Compare two firmware versions before triaging findings — see exactly which vulnerabilities were introduced or resolved between builds.

```
$ curl -s -H 'Authorization: Bearer <token>' \
'https://platform.finitestate.io/api/public/v0/versions/<baseVersionId>/compare/<targetVersionId>'
```

Handle the busy-comparison rate limit in a polling loop — respect the Retry-After header so your script backs off automatically.

```
$ curl -s -w '%{http_code}' -H 'Authorization: Bearer <token>' \
'https://platform.finitestate.io/api/public/v0/versions/<baseVersionId>/compare/<targetVersionId>' \
| grep -q '503' && echo 'COMPARISON_BUSY - retry in 2 s' &&
sleep 2
```

— Runnable endpoint with method, path and documented busy/retry behavior

[changelog-20261018-d9126399](#)

02 Conflict detection on version rescan IMPROVED

70

The `POST /api/public/v0/versions/{projectId}/rescan` endpoint returns a 409 status when a rescan is refused due to conflicting scan configs, letting callers detect and resolve the conflict before retrying rather than acting blindly on a version rename.

Detect a refused re-scan (conflicting scan configs) before your pipeline acts on a version rename.

```
$ curl -s -w '\n%{http_code}' -X POST \
  -H 'Authorization: Bearer <token>' \
  'https://platform.finitestate.io/api/public/v0/versions/<projectId>/rescan' \
  | tail -1 | grep -q '409' && echo 'SCAN_CONFIG_CONFLICT -
  resolve scan config before retrying'
```

– Named endpoint and status code, but behavior only shown via example

changelog-20261018-d9126399

THINNER COVERAGE BELOW

03 Combined Project details and permissions fetch IMPROVED

55

`GET /api/public/v0/projects` now loads Project details and permissions together in one request, eliminating the serial wait between the two fetches that previously existed.

– Named endpoint and before/after, but no runnable example given

changelog-20261018-d9126399

WAS THIS USEFUL?



DETECTION & RESPONSE

Torq

Feedback



changelog · Oct 18, 2026

CHANGELOG

RANK

Transform the enterprise SOC with Torq — combining agentic SecOps and Hyperautomation to triage, investigate, and respond to threats faster than ever.

Torq's dashboards get organizational and filtering upgrades: section separators for grouping widgets and numeric-value filters for Cases Dashboard widgets.

←→ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Numeric filtering for Cases Dashboard widgets NEW

60

Cases Dashboard widgets can now filter by numeric values — **Less Than** , **More Than** , and **Range** — to focus a widget on cases meeting a specific threshold or event count.

Focus a Cases Dashboard widget on high-volume incidents by filtering to only cases above a specific event count.

📌 In the Cases Dashboard, open a widget's filter settings, select **More Than** , and enter your event-count threshold to scope the widget to high-volume cases only.

– Names exact filter options and includes usage steps `changelog-20261018-b338ff6e`

THINNER COVERAGE BELOW

02 Section separators for dashboard widgets NEW

35

Dashboards now support separators that let you group related widgets under titled sections for cleaner reports.

– Describes benefit but no config detail or steps `changelog-20261018-b338ff6e`

WAS THIS USEFUL?



Cloudmersive

Feedback ⓘ

Private Cloud Installer v3-73 for Virus Convert Validate and NLP API (2026-10-17)

BLOG ↗

RANK

Cloudmersive is a developer API platform offering cloud-based services for file scanning and malware detection, document conversion, image processing, and data validation.

Cloudmersive's Private Cloud Installer v3-73 adds remote management support for installing updates across its Virus Scan, Convert, Validate, and NLP APIs.

↳ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 Remote update management in Private Cloud Installer IMPROVED

43

The Private Cloud Installer (v3-73) adds remote management support for Install Updates across the Virus Scan API, Convert API, Validate API, and NLP API.

– Names affected APIs but no mechanism or UI/config detail given

Private Cloud Installer v3-73 for Virus...

WAS THIS USEFUL?

👍

👎

EXPOSURE & VULNERABILITY MANAGEMENT

Qualys Patch Management

Feedback



changelog · Oct 14, 2026

BLOG ↗

RANK

Qualys Patch Management is a cloud-based security tool that automates the discovery, prioritization, deployment, and verification of software patches across an organization's IT assets.

Qualys launched TruRisk Eliminate 2.0, a new Autonomous RemOps platform built around the AI-driven TruConfidence orchestration engine and the autonomous Agent Sara, moving patch remediation from ticket-driven workflows to continuous, ring-based execution with scanner-based coverage for edge devices.

←→ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Agent Sara autonomous remediation execution NEW 70

Agent Sara is Qualys' autonomous AI agent that executes TruConfidence remediation plans in customer-defined waves and progressive rings, with automatic rollback on anomaly detection. It supports zero-touch automation for high-reliability patches on low-impact applications alongside ring-based deployment for higher-risk changes, backed by operational health monitoring and a reported sub-0.1% rollback rate.

– Explains mechanism and gives a concrete rollback-rate figure, but no command or config to invoke it.

2026-10-14-af38b70c

THINNER COVERAGE BELOW

02 TruRisk Eliminate 2.0 Autonomous RemOps platform NEW 56

TruRisk Eliminate 2.0 is described as an Autonomous RemOps platform that replaces ticket-driven patching with a continuous Decide → Execute → Prove cycle across the full attack surface.

– Describes a platform-level workflow shift but gives no setup steps. 2026-10-14-af38b70c

03 TruConfidence remediation orchestration engine NEW 51

TruConfidence is a new AI-powered Remediation and Orchestration Intelligence engine that selects the safest remediation action — patch, mitigation, uninstall, or existing control — per validated exposure.

– Names the four action types it chooses between but no usage path. 2026-10-14-af38b70c

Extends scanner-based remediation to edge and perimeter devices using existing Qualys scanners, closing coverage gaps that agent-only approaches miss.

– States the capability and rationale but no configuration detail. 2026-10-14-af38b70c

↳ ALSO FROM THESE RELEASES

.

WAS THIS USEFUL?



AGENT

Runlayer

Feedback



1.28.1480 - Slack stop button and ADK telemetry (2026-10-17)

BLOG ↗

RANK

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

Runlayer's latest release adds a Slack stop button for halting in-progress agent replies and Google ADK session telemetry for Python agents, alongside smaller usability improvements to connector search and user filters.

←→ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Google ADK session telemetry via native Python plugin NEW

70

Adds Google ADK session telemetry for Python agents via a native plugin (see Hooks SDK) that captures prompts, assistant responses, and tool activity in one conversation, with local tool enforcement and correlation to Runlayer MCP audits.

– Names mechanism and SDK but no exact command or config key

1.28.1480 - Slack stop button and ADK t...

THINNER COVERAGE BELOW

02 Slack stop button for agents NEW

50

For agents that have their own Slack app, any authorized user can now halt all in-progress replies in a thread using a stop button.

– Describes behaviour but no exact UI path or command

1.28.1480 - Slack stop button and ADK t...

03 Connector search highlighting in Manage connectors IMPROVED

50

In **Manage connectors**, search results now highlight matches, including surfacing the matching description excerpt under the connector name when the match is found in the description.

– Names exact UI location but limited mechanism detail

1.28.1480 - Slack stop button and ADK t...

04 Current user pinned first in user filters IMPROVED

45

The current user is now placed first, marked **(you)**, in user filters across Sessions, Incidents, Audit Logs, and other list views.

WAS THIS USEFUL?



GOVERN

◆ AI/LLM Security

Reality Defender

Feedback



changelog · Oct 18, 2026

CHANGELOG ↗

Reality Defender detects synthetic audio, video, images, and documents for enterprise deepfake screening.

UPDATED INTEGRATION

Social media link import – TikTok and Threads support

WAS THIS USEFUL?



OTHER / UNCATEGORIZED

Arcade

[Feedback](#)[changelog · Oct 18, 2026](#)[CHANGELOG](#)

Arcade is the actions runtime for enterprise AI agents at **PRISMA** one control point to enforce, execute, and govern every action, every time.

Arcade adds workload-identity support, single Postgres URL config, resource-quota safety gating, a Helm values schema, and credential redaction in tracing spans.

Adds **sslmode: verify-full** Postgres connection verification against the system trust store, plus support for a single Postgres connection URL as an alternative to individual config values.

Adds **deployments.resourceQuota.enabled** and **deployments.sizing.*** Helm chart values to control resource quotas and pod sizing.

Adds workload-identity support for the engine pod (**EKS IRSA**, Azure Workload Identity, **GKE** Workload Identity, or a mounted GCP key) via service account annotations, pod labels/annotations, and extra volumes/mounts.

Adds a values schema to the Helm chart that documents and type-validates every chart value.

Worker pod creation now waits for an applied resource quota before proceeding, acting as a fail-closed safety gate.

+10 more

NEW INTEGRATION[Outlook Calendar management tools](#)**UPDATED INTEGRATION**[Google Calendar rebuilt as 5.0.0 with sharing and access-inspection tools](#)**BREAKING ON UPGRADE**

- ! Google Calendar 5.0.0 is a breaking rebuild for unpinned callers — review Upgrade notes before upgrading; it also drops the **delete an event** and 'find free time slots across attendees' tools (those remain available on 4.2.4).
- ! **format=exp** on tool listings is removed.
- ! The engine image will not start on the old database schema; the migration runs automatically on **helm upgrade** but is required before the new engine image starts.
- ! Two MCP Apps feature flags are removed; the features are now always on.
- ! The engine no longer uses a curl init container to wait for worker readiness.
- ! A busybox-based init container that copied env vars into a redundant file is removed.
- ! Several images now pull from a different registry — review and update registry egress rules before upgrading.

- ! A small set of toolkit parameters that previously silently ignored `null` now reject it; callers must omit those parameters instead of passing `null`.
- ! The GitHub review-comment listing response shape changes: the `truncated` field is removed.
- ! The default curl support image chart value is deprecated and will be removed in a future chart version.

WAS THIS USEFUL?



// END OF TRANSMISSION ■

The Cyber Toolchain · curated daily · October 18, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe**

Feedback on this issue