THE CYBER TOOLCHAIN

NO. 067

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 7, 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 172 changes

Do you prefer **scroll**?

▸ // CONTENTS SHOW 29 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: GPT-5.5 | Opus 5

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

Three unrelated things stand out today: Lakera can restrict unauthorized agent tool calls at runtime; Splunk ES adds investigation

vim ? help

assistance and detection versioning; Joe Sandbox expands malware verdict depth with hundreds of new signatures, extractors, and Suricata rules.

PROTECT

Stop AI agents from calling tools they should not use

Lakera added a Dangerous Deviation detector and runtime Tool Allow/Deny List for agent tool use. Teams can move from reviewing agent behavior after the fact to enforcing which tools an agent may call while it is running.

Lakera

RESPOND

Keep investigations moving with assisted triage and versioned detections

Splunk ES added an AI Assistant for investigations, detection versioning, TAXII 2.x threat intel feed support, and analyst queue/finding group improvements. That reduces the manual context-gathering around findings and makes detection changes easier to track when an investigation depends on what logic fired.

Splunk Enterprise Security

DETECT

Extract more verdicts from malware detonations

Joe Sandbox added over 600 detection signatures, 55 malware configuration extractors, and 133 Suricata rules on Suricata 8. Analysts get more behavioral, network, and configuration detail out of a sandbox run without writing custom signatures first.

Joe Security (Joe Sandbox)

PROTECT

Turn vulnerability findings into fixes faster

Bugcrowd can now push or link Vulnerability Detections to Jira with two-way sync, Mend SCA surfaces a Top Fix in the Findings table, and Trend Vision One playbooks can remediate container image vulnerabilities automatically. The common gain is less handoff work between finding, ticketing, prioritizing, and applying the fix.

Bugcrowd · Mend SCA · Trend Vision One

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

Nmap

Feedback ⓘ changes since 2026-09-02 CODE ↗

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap gains NSE execution time and memory limits, DHX2 AFP auth, and a new `afp.uam` script argument

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-09-02
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-02
```

Adds `afp.uam` script argument to `afp.lua` to select the user authentication method, e.g. `'DHCAST128'` or `'DHX2'`

Adds DHX2 (Diffie-Hellman Exchange 2) authentication method support to `afp.lua`

Introduces execution time limits for NSE scripts: threads running more than 5 seconds without yielding are killed

Introduces memory allocation limits for NSE scripts: allocations beyond 400 MB throw an error

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-09-06 BLOG ↗ RANK

Bugcrowd is a crowdsourced security platform connecting organizations with ethical hackers to identify and report software vulnerabilities.

> Bugcrowd rolled out Jira integration for Vulnerability Detections, letting teams push or link detections—including those from Autonomous Pentest—to Jira issues with two-way sync.

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

01 **Jira linking for Vulnerability Detections** NEW 74

Vulnerability Detections can now be linked to Jira issues directly from the Bugcrowd platform, including detections surfaced by Autonomous Pentest. The Integrations panel on a Vulnerability Detection offers Push and Link actions to create a new Jira issue or connect to an existing one, with bidirectional comment sync, on-demand update pushes, and the ability to edit or remove the link at any time, keeping Jira remediation tracking in sync with risk data in Crowdcontrol.

When triaging a vulnerability detection, link it to an existing or new Jira issue to track remediation in your existing workflow.

📍 In the Bugcrowd platform, go to Customers › Linking Vulnerability Detections to Jira and follow the steps to link a Vulnerability Detection to a Jira issue.

— Names exact UI actions and sync behavior; navigation path given, no API.

Linking Vulnerability Detections to Jir…

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

StackHawk

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

StackHawk is an application security tool built for developers. Automated scanning for the entire dev pipeline combined with continuous documentation and remediation.

> StackHawk added native OS certificate trust store support for transparent TLS proxies and a non-interactive browser-based auth flow for `hawk init`.

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

01 **Native OS certificate trust store support** NEW 83

StackHawk now trusts the native OS certificate store so corporate TLS-intercepting proxies like Zscaler, Netskope, and other corporate CAs work with zero configuration. This can be disabled with the `--no-native-truststore` flag or the `USE_NATIVE_TRUSTSTORE=false` environment variable.

— Names mechanism, disable flag and env var, but no example snapshot-20260907

02 **Non-interactive browser auth for hawk init** NEW 77

A new `--browser` flag on `hawk init` enables non-interactive device-flow authentication, skipping TTY prompts and the manual key-entry menu — useful for CI pipelines or headless environments.

Authenticate in a CI pipeline or headless environment where TTY prompts would hang — the browser opens automatically for device-flow login.

```
$ hawk init --browser
```

— Exact flag with runnable command example snapshot-20260907

WAS THIS USEFUL? 👍 👎

Lakera

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Lakera is a security platform that detects and prevents AI prompt injection attacks and other LLM vulnerabilities through testing and monitoring.

> Lakera's biggest addition this window is agent-focused defense: a beta Dangerous Deviation detector and a runtime Tool Allow/Deny List to catch and restrict unauthorized agent tool calls, alongside a new early-access AI Agent Security product for agent discovery and risk assessment. It also shipped self-hosted deployment controls (`POLICY_RELOAD_INTERVAL_SECONDS` , `MAX_WARMUP_TOKENS` , `/metrics` , GPU support), expanded agentic API support (`developer` / `tool` roles, `tool_calls`), new content moderation detectors, and dashboard filtering improvements.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 Self-hosted deployment configuration and observability NEW

90

Self-hosted Guard containers gain the `POLICY_RELOAD_INTERVAL_SECONDS` environment variable to control the configurable policy reload period, and `MAX_WARMUP_TOKENS` (default 50000) to cap warmup tokens, using the minimum of `MAX_WARMUP_TOKENS` and `MAX_INPUT_TOKENS` . A new `/metrics` endpoint exposes telemetry data, Request Logging is added for monitoring solutions such as Elastic Stack, and GPU support (compatible with any CUDA 12.x GPU) accelerates request processing.

Cap warmup token usage and tune policy reload frequency when deploying a self-hosted Guard container, to reduce startup time and control reload overhead.

```ini
POLICY_RELOAD_INTERVAL_SECONDS=120
MAX_WARMUP_TOKENS=25000
```

Scrape the Guard telemetry endpoint from a Prometheus-compatible monitoring stack to track self-hosted request throughput and model latency.

```
$ curl http://<your-guard-host>/metrics
```

— Concrete env vars and endpoint with runnable examples. snapshot-20260907

02 Dangerous Deviation detector for agent tool calls NEW

83

New (beta) detector added to the Guard API and SaaS policy options that flags agent tool calls which are both dangerous and lack a warrant in trusted `user` or `system` role messages, identifying when an agent deviates outside its trusted mandate. It detects three danger types: data leakage, access and privilege escalation (sudo/admin grants, widened ACLs, public resource sharing), and system destruction (deleting or wiping persistent resources). When tool definitions are supplied in the top-level `tools` field, each detection returns an `explanation` object with `reason_code` (machine-readable, routing-safe), `reason` (user-facing text), `agent_guidance` (instruction to feed back into agent context), and `debug_info.dev_tips` (developer remediation guidance); detected events appear in Guard API logs and are counted in analytics.

— Full mechanism and field names given, but no example call shown. product docs

---

03 **Audio request debugging flag and gateway logging** NEW

65

Adds an `audio_payload` flag in Guard audio requests to expose debugging information about the audio sample, and the Gateway now logs requests to `v2/guard/audio` for monitoring purposes.

Inspect audio debug metadata when a Guard audio request is flagged, to understand why a sample was screened.

```
$ curl -X POST https://<your-guard-host>/v2/guard/audio \
  -H 'Authorization: Bearer <API_KEY>' \
  -F 'audio_payload=@sample.wav'
```

— Named flag and endpoint with a runnable example. snapshot-20260907

THINNER COVERAGE BELOW

04 **Confidence level field in Guard API breakdown** IMPROVED

55

Guard API breakdown response now includes a `result` field for each detector showing a confidence level: `l1_confident`, `l2_very_likely`, `l3_likely`, `l4_less_likely`, or `l5_unlikely`.

— Enumerates confidence levels but no usage example. snapshot-20260907

---

05 **Runtime Tool Allow/Deny List for agents** NEW

50

Adds a runtime Tool Allow/Deny List to Lakera's SaaS policy options and Defenses, letting teams control which tools an agent may call at runtime to restrict or permit specific tools and limit propagation of compromised behavior.

— Named as a policy control with UI screenshots, no config keys given. product docs

06 **User-defined prompt Allow/Deny lists** NEW 50

Adds user-defined prompt Allow/Deny lists to customise detection behaviour, with Allow/Deny list detections shown in the Request Details page.

— Names the UI location, no config syntax given. snapshot-20260907

07 **Dashboard table filtering and metadata fields** IMPROVED 50

Enhances the guardrails overview table with filtering functionality and a Policies column, adds creator and last edited by/at fields to the guardrails page, and enhances the SaaS Requests page with filtering and ordering of columns, a `request_id` field, and shareable URL-based filter views.

— Names added fields and pages but no exact steps. snapshot-20260907

08 **Developer and tool message roles in API** NEW 45

Adds support for `developer` and `tool` roles for messages in the API, and expands API coverage to include the `tool` role and `tool_calls` field, to better support agentic and tool-using workflows.

— Named fields but no example request shown. snapshot-20260907

09 **Weapons and crime content moderation detectors** NEW 45

Adds a weapons detector (firearms, bombs, missiles, bladed weapons, other personal weapons) and a crime detector as new Content Moderation detectors.

— Detailed detector scope, no configuration steps. snapshot-20260907

10 **Policy Impact Simulator in Dashboard** NEW 35

New Policy Impact Simulator in the Dashboard shows how different sensitivity levels and guardrail configurations would have affected historical traffic.

— Names Dashboard location but no walkthrough. snapshot-20260907

11 **Prompt Defense coverage for tool responses and descriptions** IMPROVED 30

Extends Prompt Defense coverage to tool responses and tool descriptions in addition to inputs, broadening protection against prompt injection surfaced through agent tool use.

— Single-line scope extension, no mechanism described. product docs

12 **AI Agent Security product (early access)** NEW 30

New AI Agent Security product, in early access, extends AI Guardrails with discovery and risk assessment for agents an organization builds and deploys.

— High-level description only, no usage details. snapshot-20260907

13 **Portuguese language support in Prompt Defense** NEW

25

Adds support for the Portuguese language in Prompt Defense.

— Single-line addition with no further detail. snapshot-20260907

14 **HTTPS endpoint for Guard** NEW

25

Adds HTTPS endpoint support for Guard.

— Bare mention, no endpoint or config detail given. snapshot-20260907

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗ RANK

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 extended Package Firewall to block malicious VS Code extensions, added incremental container scanning against a baseline, and gave AI SAST several new ways to learn organizational context — from repo files and IaC to uploaded threat models and per-finding feedback.

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

01 **Repo context file parsing for AI SAST** NEW 80

AI SAST automatically parses `AGENTS.md`, `CLAUDE.md`, skills files, and AI rules files in the repository to extract security-relevant context — such as authentication model and internal vs. external users — for every scan.

— Names exact files parsed automatically for every scan

How Endor Labs AI SAST Learns What Matt…

02 **IaC ingestion builds deployment model** NEW 80

AI SAST automatically ingests Dockerfiles, Terraform, CloudFormation, Kubernetes YAML, and Helm charts to build a deployment model used during PR scans and local developer scans, with no configuration required.

— Names exact IaC formats ingested with no setup needed

How Endor Labs AI SAST Learns What Matt…

03 **Natural-language threat model upload** NEW 75

Users can upload architecture and threat model context as plain natural language, which AI SAST verifies against the codebase and auto-classifies into project context, true-positive patterns, and false-positive patterns, scopable to a single project, a set of projects, or the entire tenant.

— Describes classification and scoping, no exact endpoint given

How Endor Labs AI SAST Learns What Matt…

04 **Package Firewall blocks malicious VS Code extensions** NEW 65

Package Firewall now blocks or warns on malicious VS Code extensions from the Microsoft Marketplace before they install, deployable via MDM tools such as Microsoft Intune, Jamf, or Kandji.

— Names deployment tools but no exact config or command product docs

05 **Incremental container scanning against baseline** NEW

60

Container scanning can now compare a new build against a baseline scan to surface only findings the build introduced, fixed, or shares with the baseline, with an option to suppress shared findings so only new ones remain.

— Explains mechanism, no exact flag or command named `product docs`

06 **Per-finding thumbs-up/down triage feedback** NEW 60

Adds per-finding thumbs-up / thumbs-down feedback during triage so teams can teach AI SAST which issue classes matter most, adjusting severity scoring for entire categories in future scans.

— Clear UI action, no exact endpoint or config given

`How Endor Labs AI SAST Learns What Matt…`

WAS THIS USEFUL? 👍 👎

Mend.io Mend SCA

Feedback ⓘ 1 RELEASE · 2026-09-06 NOTES ↗

RANK

Traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend SCA rolled out a unified vulnerability remediation experience, adding a Top Fix column to the Findings table and updating the Remediation tab for consistent fix metadata.

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

01 **Top Fix column and unified Remediation tab** NEW 63

A new Top Fix column in the Findings table surfaces the recommended remediation action per vulnerability, letting users prioritize fixes without opening individual vulnerability details. The Remediation tab was updated as part of a unified vulnerability remediation experience, standardizing fix logic and recommendations for consistent metadata across the Mend Platform.

Review the new Top Fix recommendation for each finding directly in the Findings table to prioritize remediation without opening individual vulnerability details.

📍 In the Mend Platform, go to Findings › Findings Table and locate the new Top Fix column to see the recommended fix for each vulnerability.

— Names the exact column and UI path, but mechanism of standardization is vague.

Version 26.8.2 (06-September-2026)

WAS THIS USEFUL? 👍 👎

GitHub

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

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

> GitHub Copilot adds HydraFusion, a research-preview multi-model orchestration engine selectable as a model choice, along with a CLI command to submit feedback on it.

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

01 HydraFusion multi-model orchestration in Copilot NEW

75

HydraFusion is now selectable as a model in GitHub Copilot; it dynamically orchestrates multiple models across providers using Single, Cascade, or Critique execution patterns to balance quality, cost, and latency per task. Copilot CLI supports `/feedback` to submit HydraFusion preview feedback directly from the command line without leaving the terminal.

Submit feedback on HydraFusion task quality directly from the terminal without leaving your workflow.



— Names execution patterns and a runnable CLI command, but no config or limits.

launch-20260904-1f78a8ec

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ 1 RELEASE · 2026-09-04

BLOG ↗ RANK

SonarQube is the independent AI code verification platform trusted by 7M+ developers & organizations to help build secure, reliable software with confidence.

> SonarQube Cloud now automatically analyzes four additional languages — MuleSoft DataWeave, Gosu, Groovy, and PowerShell — with no plugins or manual configuration required.

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

01 Automatic analysis for DataWeave, Gosu, Groovy, and PowerShell

NEW 80

SonarQube Cloud adds automatic, no-plugin-required analysis for four new languages: `MuleSoft DataWeave` integration flows for bugs, security vulnerabilities, and maintainability issues across GitHub, GitLab, Bitbucket, and Azure DevOps; `Gosu` for deterministic bug, vulnerability, and maintainability analysis on Guidewire-based insurance and enterprise systems; `Groovy` covering Jenkins pipeline definitions, Gradle build scripts, and standalone Groovy applications with a dedicated rule set; and `PowerShell` infrastructure automation scripts, detecting hardcoded credentials, injection risks, unsafe command execution, incorrect logic, and unsafe type handling.

Connect a Groovy-heavy Jenkins or Gradle repository to SonarQube Cloud so pipeline definitions and build scripts are analyzed automatically on every pull request.

> 📍 In SonarQube Cloud, go to your organization › `Add a project`, select your GitHub/GitLab/Bitbucket/Azure DevOps repository containing Groovy files (Jenkinsfiles, build.gradle, etc.), and run your first analysis. Groovy is picked up automatically — no plugin installation or language configuration required.

— Names four languages with concrete risk types and platforms, plus a setup example

launch-20260904-f53d30da

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

1 RELEASE · 2026-08-17 BLOG ↗ RANK

GitGuardian Detection Engine scans code repositories to identify and alert on exposed secrets, credentials, and sensitive data.

> GitGuardian Detection Engine 2.170 expands secret coverage with six new detectors and a new token analyzer, while sharpening existing detectors for Docker Hub, New Relic, Fastly, and Scaleway.

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

01 **False positive reductions for Fastly and Scaleway tokens**

IMPROVED 65

Reduces false positives for `fastly_personal_token` by rejecting proxy-subscription node passwords and 32-char prefixes of longer base64url values. Reduces false positives for `scaleway_token` by rejecting node UUIDs from proxy-subscription configs.

— Specific detector names and exact rejection logic described

Detection Engine Updates Version 2.170 …

THINNER COVERAGE BELOW

02 **Docker Hub personal access token detection** IMPROVED

55

Extends the Docker Credentials detector to catch Docker Hub personal access tokens used with `docker login --username` / `-u`.

— Names exact CLI flags but no further mechanism Detection Engine Updates Version 2.170 …

03 **Revoked key detection for New Relic checkers** IMPROVED

55

New Relic Key, New Relic Insights Insert Key, and New Relic APM License Key checkers can now identify revoked keys via an empty POST request.

— Explains mechanism but no endpoint or config given

Detection Engine Updates Version 2.170 …

04 **New detectors for Firebase, Coze, Keycloak, Fish Audio, Azure Custom Vision, Apify**

NEW 51

> Adds new detectors for Firebase ID tokens, Coze access tokens, Keycloak access tokens, Keycloak refresh and offline tokens, Fish Audio API keys, and Azure Custom Vision Training keys, plus a new analyzer for Apify API tokens.
>
> — Names six detectors and one analyzer but gives no mechanism

Detection Engine Updates Version 2.170 …

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Radare2

Feedback ⓘ 1 RELEASE · 2026-09-06 NOTES ↗ CODE ↗ RANK

UNIX-like reverse engineering framework and command-line toolset

> Radare2 6.2.2 adds a JNI analysis plugin, initial NVIDIA PTX SASS support, a native Haiku debugger, and deeper .NET/CLR, ELF relocation and calling-convention analysis.

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

01 **Calling convention and variable recovery improvements**

IMPROVED 55

Adds `dyncc` floating-point argument registers to calling convention analysis, arm64 Swift calling convention support for FPU arguments, and handling of subregister arguments on x64 during variable recovery.

— Names three specific argument-handling improvements, no example 6.2.2

02 **JNI and Java descriptor analysis plugin** NEW 50

Adds the `jni` analysis plugin to resolve typed JNI table calls, interface tables, indirection, and materialize JNI native methods during analysis, along with structured Java and JNI descriptor parsing.

— Names the plugin but no usage example or command shown 6.2.2

03 **Native Haiku debugger** NEW 50

Adds native debugger support for Haiku (x86-32/64), including data watchpoints and asynchronous thread exit handling.

— Specifies architectures and two capabilities, no invocation shown 6.2.2

04 **.NET/CLR analysis support** NEW 50

Adds CIL token resolution and stack ESIL for .NET/CLR binaries, plus parsing of CLR exception clauses in .NET method bodies.

— Groups two named .NET/CLR capabilities, no usage detail 6.2.2

05 **Mach-O dSYM sideloading** NEW 50

Adds sideloading of Mach-O dSYM files matched by architecture slice and UUID.

— Explains matching mechanism (slice and UUID), no command shown 6.2.2

06 **ELF relocation handling improvements** IMPROVED 40

Adds support for Android-specific ELF relocations and full expansion of ELF RELR bitmap relocations.

— Names two relocation types but no mechanism detail 6.2.2

07 **ELF eh_frame Itanium exception parsing** NEW 40

Adds parsing of Itanium exception regions from ELF `eh_frame`.

— Names the exact ELF section but no usage detail 6.2.2

08 **RBinSymbol.hsize header padding field** NEW 40

Adds the `RBinSymbol.hsize` field for header metadata paddings.

— Names exact struct field, no example of use 6.2.2

09 **NVIDIA PTX SASS architecture support** NEW 35

Adds initial encode/decode/pseudo support for the NVIDIA PTX SASS architecture.

— Names the architecture but 'initial' scope undefined 6.2.2

10 **s390 pseudo-code parser** NEW 30

Adds an `s390` pseudo-code parser.

— Bare name with no mechanism described 6.2.2

11 **pcap stream selection support** NEW 25

Adds pcap stream selection support.

— Bare mention with no mechanism or command 6.2.2

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

Joe Security (Joe Sandbox)

Feedback ⓘ

1 RELEASE · 2026-09-07 BLOG ↗ RANK

Automated Malware Analysis - Development and Licensing of Automated Malware Analysis Tools to Fight Malware

> Joe Sandbox v45 adds over 600 new detection signatures (50 behavior signatures and 575 YARA rules), 55 malware configuration extractors, and 133 Suricata rules on an upgraded Suricata 8 engine, alongside improvements to DOM/phishing detection, QR code analysis, Live Interaction responsiveness, and anti-evasion protections including new macOS shutdown protection.

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

01 **New malware configuration extractors for 10 families** NEW

65

Adds 55 new malware configuration extractors — including `ZigClipper`, `BlakcSee Stealer`, Sirius RAT, `OverLordRat`, `CastleRat`, `WikiKit`, `PhantomGate`, Remus Stealer, `SantaStealer`, and `GenericBot` — to automatically extract IOCs and family configuration data without manual reverse engineering.

— Names specific extractors and their automated purpose

Joe Sandbox v45 - Green Opal (2026-09-0…

02 **New Suricata rules and Suricata 8 engine upgrade** NEW

65

Adds 133 new Suricata rules (30 anomaly, 88 malware, 15 phishing) covering CastleRAT, LummaC, LxBase RAT, njRAT, Phoenix Stealer, SilentNet, ErrTraffic v2 ClickFix, EvilTokens, Tycoon2FA, and WikiKit/Sneaky2FA, and updates the underlying Suricata engine to Suricata 8.

— Names threats covered, rule counts, and engine version

Joe Sandbox v45 - Green Opal (2026-09-0…

THINNER COVERAGE BELOW

03 **New behavior signatures for emerging malware families** NEW

55

Adds 50 new behavior signatures covering emerging malware families including RatonRAT, REMUS Stealer, EtherHiding, GehennaLocker, PureLogs Stealer, PhantomGate, QuimaRAT, and others.

— Names families detected but no mechanism beyond signature count

Joe Sandbox v45 - Green Opal (2026-09-0…

04 **New YARA rules for ransomware, stealers, RATs and phishing**

NEW

35

Adds 575 new custom YARA rules to detect ransomware, stealers, RATs, and phishing frameworks.

— Only category names given, no specific families or usage path

Joe Sandbox v45 - Green Opal (2026-09-0…

05 **macOS shutdown evasion protection** NEW

35

Adds macOS shutdown protection to prevent analysis evasion via system shutdown on macOS samples.

— Names platform and evasion vector but no mechanism

Joe Sandbox v45 - Green Opal (2026-09-0…

06 **Improved DOM extraction and phishing detection** IMPROVED

30

Improves DOM extraction and phishing detection for better visibility into dynamic web content, obfuscated scripts, and redirect-based phishing pages.

— Describes scope but no concrete mechanism or metric

Joe Sandbox v45 - Green Opal (2026-09-0…

07 **Improved QR code reader analysis** IMPROVED

30

Improves QR code reader analysis to uncover embedded URLs and data in QR-based phishing samples and documents.

— States purpose but no before/after detail Joe Sandbox v45 - Green Opal (2026-09-0…

08 **Improved Live Interaction responsiveness** IMPROVED

20

Improves Live Interaction responsiveness and performance for smoother manual analyst interaction with analysis machines.

— Generic performance claim with no specifics Joe Sandbox v45 - Green Opal (2026-09-0…

09 **Improved VM and analysis-system detection evasion prevention**

IMPROVED

20

Improves prevention of VM and analysis system detection techniques used by evasive malware.

— Vague statement with no technique names or metrics

Joe Sandbox v45 - Green Opal (2026-09-0…

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Mobile Verification Toolkit

Feedback ⓘ

1 RELEASE · 2026-09-07 NOTES ↗ CODE ↗ RANK

MVT (Mobile Verification Toolkit) helps with conducting forensics of mobile devices in order to find signs of a potential compromise.

> MVT speeds up iOS backup decryption with parallel workers, adds a sysdiagnose IOC-scanning module, and fixes data-loss bugs in both iOS backup decryption and Android settings parsing.

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

01 **Concurrent workers for iOS backup decryption** NEW 90

Adds a `--jobs` option to `mvt-ios decrypt-backup` to control the number of concurrent decryption workers, accepting an integer from 1 up to `MAX_DECRYPT_WORKERS` and defaulting to `DEFAULT_DECRYPT_WORKERS`.

Decrypt a large encrypted iOS backup faster by parallelising file decryption across multiple workers.

```
$ mvt-ios decrypt-backup --jobs 8 -p <password> -d /tmp/decrypted
/path/to/backup
```

— Exact flag, bounds, and a runnable command are given v2026.9.7

THINNER COVERAGE BELOW

02 **iOS backup decryption dependency fix** IMPROVED 55

Replaces the unmaintained `iOSbackup` dependency with `iphone_backup_decrypt` for iOS backup decryption, fixing decrypted-file truncation that was caused by inaccurate sizes in backup metadata.

— Clear before/after fix but nothing for reader to invoke v2026.9.7

03 **SysdiagnoseInfo module for check-sysdiagnose** NEW 50

Adds a `SysdiagnoseInfo` module to `check-sysdiagnose`, extending IOC scanning to sysdiagnose archives.

— Names the module but no usage detail given v2026.9.7

04 **Multi-line dumpsys settings parsing** IMPROVED 50

Parses `dumpsys settings` output as per-record results in Android analysis, capturing multi-line values and metadata that the previous per-line regex truncated or dropped.

— Explains behaviour change but no command or path shown v2026.9.7

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

nfdump

Feedback ⓘ 1 RELEASE · 2026-09-06 NOTES ↗ CODE ↗ RANK

Netflow processing tools

> nfdump 1.7.9 adds geo-enrichment for exporter/router tags, IPFIX nanosecond timestamp support, new ASN aggregation fields, and restructured non-root Docker builds.

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

01 **srcasn/dstasn/ext aggregation fields** IMPROVED 75

Adds the missing `srcasn`, `dstasn`, and `ext` fields to flow aggregation, allowing per-AS traffic summaries such as `nfdump -r /flow_base_dir/router1/nfcapd.202501011200 -A srcasn,dstasn`.

Aggregate flows by source and destination ASN to produce per-AS traffic summaries

```
$ nfdump -r /flow_base_dir/router1/nfcapd.202501011200 -A srcasn,dstasn
```

— Runnable command with named aggregation fields v1.7.9

THINNER COVERAGE BELOW

02 **Geo-enriched exporter/router address tag** NEW 50

Adds a `%gra` geo-enriched router/exporter address tag for use in `fmt` output, and includes geo info in `output_raw` output.

— Names exact format tag but no further mechanism detail v1.7.9

03 **Non-root multi-stage Docker builds** IMPROVED 50

Docker builds now use a multi-stage process producing separate `nfcapd` (collector) and `nfdump` (client) targets, which run as non-root by default.

— Describes build change but no exact commands given v1.7.9

04 **IPFIX nanosecond time elements 156/157** NEW 45

Adds support for IPFIX nanosecond-precision time elements 156 and 157, compatible with VPP 17.07.

— Names elements and compatibility but no usage steps v1.7.9

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Hunters Pathfinder AI

Feedback ⓘ

1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

> Hunters Pathfinder AI ships a unified anomalous-login detector that consolidates and deprecates twelve legacy login/EDR detectors, adds Windows password-change enrichment with a companion risk-scoring model, and rolls out four new AXON threat hunting reports covering encoded command execution, Azure SPN sign-in anomalies, suspicious authentication tooling, and public cloud storage exposure.

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

01 Unified anomalous login detector replaces twelve legacy detectors

80

BREAKING

Adds the Context Unified Anomalous Login detector, which evaluates successful anomalous logins across identity providers, SaaS applications, cloud platforms, VPNs, and other integrated login sources using per-user behavioral baselines covering geolocation, IP address, and user agent category. It replaces twelve legacy anomalous-login and EDR-presence detectors — `CloudTrail anomalous console login`, Anomalous login by an Azure account, Office365 anomalous login, Okta anomalous login, Okta successful login from host without an EDR agent, Okta successful administrative access from host without an EDR agent, Successful Duo authentication from host without an EDR agent, Successful Azure portal sign-in from machine without EDR agent, Successful Azure CLI sign-in from machine without EDR agent, Successful AWS console login from host without an EDR agent, Successful G Suite login from computer without EDR agent, and G Suite successful authorization of token to GCP related interface without EDR agent — all of which are deprecated and scheduled for removal on July 6, 2026.

— Names replaced detectors and removal date; no migration steps given snapshot-20260907

02 Windows password-change enrichment and scoring model NEW

80

Adds a Windows Recent Password Change enrichment that queries Windows Security Event Logs for Event ID 4723 (self-initiated password change) and Event ID 4724 (privileged password reset), surfacing last change time, initiating user, and source computer as lead context. A companion Windows Recent Password Change Model scoring layer applies a configurable confidence modifier to reduce noise from password-reset-triggered alerts and increase confidence when a password change may indicate account takeover.

— Event IDs and modifier logic named, but no config steps shown snapshot-20260907

03 **Four new AXON Threat Hunting Reports** NEW 80

Adds a report for multi-layer Base64-encoded command execution covering defense evasion and obfuscation across Windows, Linux, and macOS process creation events; a report for Azure Service Principal sign-ins via anomalous user agents using a 45-day behavioral baseline with automatic suppression when a corresponding app registration update occurred within 24 hours of the sign-in; a `Suspicious User-Agent Observed in Azure Non-Interactive/SPN Sign In Activity` report identifying authentication activity from offensive-related distributions and tools including Kali Linux, Parrot Linux, BackBox Linux, BlackArch Linux, Scout Suite, TrailBlazer, and TruffleHog; and a report for publicly accessible cloud storage resources across AWS S3, Azure Blob Storage, and Google Cloud Storage identifying externally reachable buckets and containers with anonymous enumeration or access. Findings from all four campaigns can be reviewed by navigating to Axon Reports in the Hunters platform.

Review AXON threat hunting findings for multi-layer encoded command execution, anomalous SPN sign-ins, and cloud storage exposure after the new reports land in your tenant.

> 📍 In the Hunters platform, navigate to Axon Reports to view findings from the Multi-Layer Encoded Command Execution, Azure Service Principal Sign-In via an Anomalous User Agent, Suspicious User-Agent in Azure Non-Interactive/SPN Sign-In Activity, and Publicly Accessible Storage Resources campaigns.

— Each report named with mechanism and a UI path to review findings snapshot-20260907

└──▷ **BREAKING ON UPGRADE**

! The following detectors are deprecated and will be removed on July 6, 2026, replaced by Context Unified Anomalous Login : `CloudTrail anomalous console login` , Anomalous login by an Azure account , Office365 anomalous login , Okta anomalous login , Okta successful login from host without an EDR agent , Okta successful administrative access from host without an EDR agent , Successful Duo authentication from host without an EDR agent , Successful Azure portal sign-in from machine without EDR agent , Successful Azure CLI sign-in from machine without EDR agent , Successful AWS console login from host without an EDR agent , Successful G Suite login from computer without EDR agent , G Suite successful authorization of token to GCP related interface without EDR agent .

WAS THIS USEFUL? 👍 👎

Cisco / Splunk Splunk Enterprise Security

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Splunk Enterprise Security is a platform that collects, analyzes, and correlates security data to detect threats and respond to incidents.

> Splunk ES 8.2 introduces an AI Assistant for investigations, TAXII 2.x threat intel feed support, expanded REST APIs, detection versioning, KV Store retention policies, and analyst queue/finding group improvements, while replacing the Investigation Workbench UI with Mission Control in a breaking, one-way upgrade.

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

01 Analyst queue filtering, notes and refresh controls IMPROVED

85

The analyst queue gains filter buttons — Findings, Investigations, Finding groups, and All types — for quick type-based filtering, configurable note and note title requirements (optional or required) when analysts update a finding or investigation, an option to hide duplicate findings already nested under an investigation, ability to sync changes made in an investigation or finding group to all included findings, ability to view notes on findings or finding groups included in an investigation, and manual and auto-refresh controls now available in the quick actions menu. Findings that are part of an investigation are now hidden from the top-level analyst queue by default, and search ID reuse improves load times across the analyst queue and investigations by reusing cached search job results.

— Names all filter labels and behaviors, UI-based actionability snapshot-20260907

02 Mission Control replaces Investigation Workbench BREAKING

80

The Investigation bar, Investigation Workbench, and Investigation dashboard are removed and replaced by the Mission Control UI. The `edit_timeline` and `manage_all_investigations` capabilities have been removed, along with workbench-related views (`ess_investigation_list`, `ess_investigation_overview`, `ess_investigation`), the investigation type macro configuration, sequence templates, enhanced workflows, Incident Review row expansion, and the Workbench Profile, Workbench Panel, and Workbench Tab content types on the Content management page. The ability to require a note when an analyst updates a finding in the analyst queue, previously available in 7.3, is no longer supported in 8.x.

— Enumerates every removed view and capability by name snapshot-20260907

03 **Detection version activity, audit trail and editor testing**

NEW

75

Adds a detection version activity view to track the history of changes to a detection, plus a detection audit trail showing when detections are turned on or off, modified, or deployed, including who made the change and when. The detection editor now supports testing detections and reviewing search results directly in the editor, along with SPL query validation for custom finding-based detections.

— Describes mechanism and editor location, no exact command snapshot-20260907

04 **Breaking upgrade path to 8.x** BREAKING

75

Upgrading to Splunk Enterprise Security 8.x is a one-way operation — you can no longer access any investigations created prior to the upgrade. Upgrades from versions 6.x and earlier are not supported; you must first upgrade to version 7.3.2 before upgrading to 8.x, and you cannot upload Splunk Enterprise Security 8.x on an on-premises Splunk Enterprise 10.x deployment via the UI — installation must be done from the command line.

— Gives exact minimum version and CLI-only install path snapshot-20260907

05 **Expanded REST APIs for findings, investigations and notes**

NEW

70

Expanded REST APIs let you create findings, add findings to investigations, and create, read, update, and delete notes; these new Splunk Enterprise Security APIs are now searchable using the SPL `REST` command.

— Names SPL REST command but no endpoint paths given snapshot-20260907

06 **AI Assistant for investigations** NEW

65

New AI Assistant for investigations can summarize findings, generate SPL searches, and produce investigation reports, including summarizing individual findings directly in the analyst queue. Analysts can choose between Frontier or Splunk-hosted models to meet compliance requirements.

— Names model choice but no exact command or config path snapshot-20260907

07 **KV Store retention policy** NEW

60

Adds a KV Store retention policy that automatically removes old records from KV Store collections based on a configured time-based or size-based policy.

— Explains mechanism but no config key named snapshot-20260907

08 **Lookback and overlap finding groups with reopen control** NEW

60

Adds configurable behavior for whether closed finding groups are reopened when additional findings are added, along with lookback finding groups that group historical findings based on the first time a detection runs, and overlap finding groups that catch edge cases that might represent risk.

— Names three grouped behaviors, no config path `snapshot-20260907`

THINNER COVERAGE BELOW

09 **TAXII 2.0/2.1 threat intel feeds** NEW 53

Adds support for TAXII version 2.0 and TAXII version 2.1 threat intelligence feeds, configurable via Add a TAXII 2 feed .

— Names versions and UI entry point, no further mechanism `snapshot-20260907`

10 **Hybrid Cloud/SOAR pairing support** NEW 40

Adds hybrid pairing support between Enterprise Security (Cloud) and a single Splunk SOAR (On-premises) instance.

— Bare capability statement, no setup detail `snapshot-20260907`

└──▷ **BREAKING ON UPGRADE**

! Upgrading to Splunk Enterprise Security 8.x is a one-way operation — you can no longer access any investigations created prior to the upgrade.

! Upgrades from versions 6.x and earlier are not supported; you must first upgrade to version 7.3.2 before upgrading to 8.x.

! You cannot upload Splunk Enterprise Security 8.x on an on-premises Splunk Enterprise 10.x deployment via the UI — installation must be done from the command line.

! The Investigation bar, Investigation Workbench, and Investigation dashboard are removed and replaced by the Mission Control UI.

! The `edit_timeline` and `manage_all_investigations` capabilities have been removed.

! Workbench and workbench-related views (`ess_investigation_list` , `ess_investigation_overview` , `ess_investigation`) have been removed.

! Configuring the investigation type macro is no longer available.

! Sequence templates are no longer available.

! Enhanced workflows are no longer available.

! Incident Review row expansion is no longer available.

! The Content management page no longer includes Workbench Profile , Workbench Panel , or Workbench Tab content types.

! The ability to require a note when an analyst updates a finding in the analyst queue (previously available in 7.3) is no longer supported in 8.x.

WAS THIS USEFUL? 👍 👎

Databricks Panther

Feedback ⓘ 5 RELEASES · SEEN 2026-09-07

NOTES ↗ RANK

Panther closes the loop across the SOC with AI that investigates alerts, improves detections, and reduces noise over time.

> Panther's biggest additions this window are a new Log Forwarder for shipping on-prem logs (including Palo Alto NGFW) that's now manageable via REST API and Terraform, plus expanded AI-driven triage with auto-triage deferral and an on-demand Slack Bot AI Triage button. It also picked up an Infrastructure Cost Dashboard, a batch of new log sources and schema expansions, and several settings/permissions and UI refinements.

└──▷ WHAT SHIPPED · 16 FEATURES most completely described first ? what's the number?

01 **Log Forwarder for on-prem log ingestion** NEW 75

Adds a new Panther Log Forwarder source for sending on-prem logs to Panther, alongside a new log source for Palo Alto Next-Generation Firewall (NGFW) logs to bring on-prem firewall activity into Panther. Log Forwarder was then added as a transport option specifically for ingesting Palo Alto NGFW logs, and Log Forwarder sources can now be managed as code via the REST API and Terraform provider.

Ingest on-prem Palo Alto NGFW syslog traffic into Panther without routing it through a cloud forwarder first.

> 📍 In Panther, go to Log Sources › Create New Source, select Palo Alto Next-Generation Firewall, configure the listener or Log Forwarder endpoint, and point your NGFW syslog output at the provided address.

Route Palo Alto NGFW logs into Panther using the new Log Forwarder transport instead of direct S3 or SQS ingestion.

> 📍 In Panther, go to Log Sources › Add Log Source, select Palo Alto Next-Generation Firewall, and choose Log Forwarder as the transport option.

— Named source and transport evolution, but no exact Terraform resource names v1.128 v1.127 v1.126

02 **AI scheduled prompt permissions and code management** NEW 70

Adds `ManageAIResponses` permission so admins can manage other users' private scheduled prompts, and `ViewAIPrivateResponses` permission granting view-only

admins read access to private scheduled prompts without modify rights; scheduled prompts can now also be managed as code via GitHub PR.

Grant a view-only admin visibility into private AI scheduled prompts without allowing modifications.

> 📍 In Panther, go to Settings › Users Management, select the user, and assign the `ViewAIPrivateResponses` role.

— Named permissions and UI path but no API endpoint v1.127

---

03 **Time-bounded correlation rule matching** NEW 70

Adds `WithinTimeFrameMinutes` option to group correlation rules for CTE-based matching, enabling time-bounded correlation across grouped events.

Enable time-bounded correlation across grouped events — useful for detecting brute-force patterns that must occur within a rolling window.

```yaml
WithinTimeFrameMinutes: 10
```

— Config key and example shown, mechanism explained clearly v1.128

---

04 **AI-powered alert triage automation and Slack Bot** NEW 65

Adds AI auto-triage capability that can defer alert notifications until triage completes. The Panther Slack Bot gained an AI Triage button to trigger triage on demand plus `@mention` support for follow-up questions in Slack threads (requires updating your Slack app manifest), and Panther AI Slack messages now render full markdown (bold, links, tables, and lists).

Trigger on-demand AI triage from Slack when an alert fires, without leaving the thread.

> 📍 In the Panther Slack Bot alert thread: 1. Click the AI Triage button to start automated triage. 2. After triage completes, @mention the Panther bot in the thread to ask follow-up questions. 3. If the button or @mention is missing, update your Slack app manifest as described in the Panther docs.

— Concrete Slack steps given but no API or config details v1.128 v1.127

THINNER COVERAGE BELOW

---

05 **New log sources and schema expansions** NEW 50

Adds ingestion of Cursor audit logs over an HTTP webhook; adds Claude Code and Claude Cowork as new log sources; adds `Tracebit.AlertSummary` log type and

extends `Tracebit.Alert` with Google Cloud principals; adds `Material.IssueChange` and `Material.AuditLog` log types; adds Identity Threat Protection fields to the `Okta.SystemLog` schema; and expands the Google Workspace schema with additional Activity types and new Reports API fields.

Stream Cursor audit logs into Panther for SaaS activity monitoring by configuring an HTTP webhook source.

> 📍 In Panther, go to Log Sources › Create New Source, select HTTP Webhook, choose Cursor as the log type, and copy the generated webhook URL into your Cursor audit log export settings.

— Multiple named schemas listed but little mechanism per item v1.128 v1.127 v1.126

---

06 **Force sign-out of all user sessions** NEW 50

Enables signing a user out of all active sessions from the Users Management settings page.

Force a compromised or departing user out of all active Panther sessions immediately from the admin console.

> 📍 In Panther, go to Settings › Users Management, select the target user, and choose Sign out of all active sessions.

— Clear UI path provided, named settings page v1.127

---

07 **Archived schema filtering** IMPROVED 45

Schemas page now hides archived schemas by default, with a Status filter to show them.

— Named filter field but no exact UI path v1.126

---

08 **G Suite Drive owner email normalization** IMPROVED 45

Normalizes G Suite Drive document-owner email into `p_any_emails` for cross-source correlation.

— Named field but no usage example given v1.127

---

09 **Infrastructure Cost Dashboard** NEW 40

Adds an Infrastructure Cost Dashboard that tracks Panther usage and billing trends, letting admins review current spend and forecast costs or identify unexpected usage spikes.

Review current Panther infrastructure spend and billing trends to forecast costs or identify unexpected usage spikes.

> 📍 In the Panther console, navigate to the Infrastructure Cost Dashboard to view usage metrics and billing trends.

– UI dashboard named, no metrics or limits given `snapshot-20260907`

---

10 **GitHub PR workflow for config-as-code changes** IMPROVED

40

Enables creation of GitHub PRs for data models, global helpers, and deletions, and later adds support for deleting Data Models, Global Helpers, and Queries via a GitHub PR.

– Named entities but no concrete steps or endpoint given `v1.127` `v1.126`

---

11 **Updated left navigation sidebar** IMPROVED 35

Updated left navigation brings a flatter structure, with settings and the user menu now built into the sidebar.

– UI description only, no exact navigation steps `v1.128`

---

12 **Redesigned Settings page** IMPROVED 35

Redesigned Settings page consolidates all configuration into a single full-page view with a persistent sidebar.

– UI description only, no navigation path given `v1.126`

---

13 **Databricks-backed scheduled lookups** NEW 30

Scheduled lookups now support Databricks-backed queries.

– Named integration but no usage example or config shown `v1.129`

---

14 **STIX/TAXII log enrichment** NEW 30

Adds STIX/TAXII enrichment for ingested logs.

– Bare mention, no mechanism or usage details `v1.128`

---

15 **Actionable enrichment upload errors** IMPROVED 30

Enrichment and custom enrichment uploads now surface actionable errors for duplicate keys and invalid files.

– Describes behavior, no specifics on error format `v1.126`

---

16 **Enriched-field filtering for normalized events** IMPROVED

25

Normalized event filters can now filter events using Panther-enriched fields.

– One-line description, no example or mechanism `v1.129`

WAS THIS USEFUL?  

Databricks

Feedback

16 RELEASES · 2026-09-01 → 2026-09-04

BLOG ↗

RANK

Databricks is a unified data and AI platform that combines data engineering, analytics, and machine learning on Apache Spark infrastructure.

> Databricks shipped a broad round of dashboard upgrades (embedded filter APIs, cross-dataset filtering, Slack/Teams data attachments, and more), took Genie Code scheduled tasks to General Availability, added three new hosted foundation models (GPT-6 Astra, Gemini 3.8 Flash, Claude Fable 5.1), and introduced Beta features including Genie One web search, OpenSharing for metric views, Git Folder Serverless, and AI Search support for struct/map columns and filter-only queries.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 Genie Code scheduled tasks reach General Availability NEW

80

Genie Code scheduled tasks are now generally available: a prompt can be scheduled to run on a recurring schedule via chat prompt or the Schedules menu, producing a continuable Genie chat with results for each run. Scheduled tasks will soon be enabled by default for workspaces with the compliance security profile enabled.

Set up a recurring analysis task when you want Genie Code to automatically run a prompt (e.g. a weekly data quality report) without manual intervention.

📍 In the Genie Code chat, type your prompt (e.g. 'Summarize data quality issues in the sales table') and then go to Schedules › Create schedule to set the recurrence. Each run produces a new continuable Genie chat with the results.

– Includes exact UI steps and a runnable example prompt.

Genie Code scheduled tasks are now gene… | Genie Code scheduled tasks will soon be…

02 Cross-dataset filtering through relationships NEW

80

Dashboards add cross-dataset filtering through relationships (Public Preview): filters, selections, and parameters applied to one dataset automatically carry across to widgets built on related datasets, configured via Data > Relationships in the dashboard editor.

Enable cross-dataset filtering so that a region filter on a sales dataset automatically narrows widgets built on the related returns dataset.

📍 In the dashboard editor, open Data › Relationships, click `Add relationship`, select the two datasets and the shared key field, then save. Filters applied to either dataset will now propagate to the other.

— Exact UI navigation and preview status given with a worked example.

Dashboard enhancements and bug fixes (2…

03 Data attachments in dashboard subscriptions NEW 75

Dashboard subscriptions delivered to Slack and Microsoft Teams can now include `Data attachments` in CSV, TSV, or Excel format, configured via Subscriptions > Add subscription in the dashboard.

Deliver a scheduled dashboard snapshot to a Teams channel with a CSV data attachment so recipients can slice the raw numbers without opening Databricks.

📍 In the dashboard, open Subscriptions › Add subscription, set the destination to your Microsoft Teams channel, and under `Data attachments` select CSV. Save and schedule the subscription.

— Names exact formats and the UI steps to configure it.

Dashboard enhancements and bug fixes (2…

04 New foundation models on Unity Gateway and Model Serving NEW 70

Databricks added three new hosted models accessible via Foundation Model APIs: `OpenAI GPT-6 Astra` on Unity Gateway (pay-per-token, supporting enterprise reasoning, structured document processing, coding, agentic search, multimodal inputs, and long-context workloads), `Google Gemini 3.8 Flash` in Model Serving, and `Anthropic Claude Fable 5.1` on Unity Gateway.

— Names three specific models and their access surfaces, no exact call shown.

OpenAI GPT-6 Astra is now available on … Google Gemini 3.8 Flash now available a… Anthropic Claude Fable 5.1 now availabl…

05 AI Search struct/map columns and filter-only queries NEW 60

AI Search now supports indexing and filtering of struct and map columns, and supports filter-only queries that return matching rows without requiring vector or keyword retrieval.

— Describes new query and indexing behavior without a usage example.

AI Search supports struct and map colum…

06 Git Folder Serverless shared environment (Beta) NEW

60

Git Folder Serverless (Beta) supports a `pyproject.toml` `-managed` environment shared across all notebooks and files within a Git folder.

— Names the exact config file but omits setup steps.

Git Folder Serverless is in Beta (2026-…

THINNER COVERAGE BELOW

07 **Ranked table calculations in /importBI Genie Code command**

IMPROVED

55

The `/importBI` Genie Code command, used to import reports from Power BI and Tableau, now supports ranked table calculations.

— Names the exact command but gives no worked example.

Dashboard enhancements and bug fixes (2…

08 **Genie access to OpenAI-served models on Databricks** NEW

55

Genie One and Genie Agents can now use models served through OpenAI on Databricks when the `partner-powered AI features` setting is enabled.

— Names the enabling setting but not its exact location.

Genie One and Genie Agents can now use …

09 **Dashboard theme color mappings** NEW

50

Dashboard theme settings add color mappings, letting you assign a specific color to a named value that is used consistently across all visualizations, both when coloring by field and coloring by series.

— Describes behavior but no config key or navigation given.

Dashboard enhancements and bug fixes (2…

10 **Additional dashboard UX refinements** IMPROVED

50

Text widgets, widget titles, and descriptions now support rich text formatting for dynamic data references; dashboard variables support multiple-field selection so viewers can add, remove, and swap several fields at once instead of one at a time; dashboard variable option lists show their source dataset as a subtitle; downloads from a metric view now use the display names defined for its fields; scrollbars remain always visible on scrollable faceted charts; and Genie Code automatically places new widgets in available layout space when adding them to a dashboard.

— Several minor UX tweaks with no config or command names.

Dashboard enhancements and bug fixes (2…

11 **Web search in Genie One (Beta)** NEW 50

Genie One can now search the public web to answer questions requiring current information — such as recent releases, third-party documentation, or news — automatically deciding when to search and citing each source as an openable link in its responses.

— Describes behavior and citations but no UI path or example.

Web search in Genie One (Beta) (2026-09… | Genie One web search is in Beta (2026-0…

12 **Genie One mobile notifications and embedded dashboard UX** NEW 50

The Genie mobile app now sends push notifications for scheduled tasks to iOS devices (Public Preview); embedded dashboards now hide the Ask Genie button when the feature is not supported.

— States platform and preview status but no setup steps.

Genie One enhancements (2026-09-03)

13 **Embedded dashboard filter APIs** NEW 45

Databricks dashboards add embedded filter APIs that let host applications programmatically set and retrieve filter values in embedded dashboards.

— Describes capability without naming the endpoint or method.

Dashboard enhancements and bug fixes (2…

14 **OpenSharing for metric views (Beta)** NEW 45

Metric views can now be shared with Databricks users in other metastores or accounts using OpenSharing (Beta).

— States the capability without setup steps or endpoint.

Share metric views with OpenSharing (Be…

15 **Apache Arrow Flight support for Zerobus Ingest (GA)** IMPROVED 45

Zerobus Ingest in Lakeflow Connect now generally supports Apache Arrow Flight.

— Names the feature and GA status, no further mechanism.

Apache Arrow support for Zerobus Ingest…

16 **Environment version 6 for compute** NEW 45

Environment version 6 is now available as a selectable base environment for serverless and standard classic compute.

— Names the version and compute types but no selection steps.

Environment version 6 is now available …

17 **Genie Agents tool descriptions and model upgrade** IMPROVED

35

Genie Agents now displays the function description when a SQL function is attached as a tool, giving users immediate context on what each function does; the underlying large language models used in Genie Agents chat mode have also been upgraded.

— No specific model names given for the LLM upgrade.

Genie Agents enhancements (2026-09-03)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI shipped FIPS 140-3 compliant deployment images, inline entity creation inside investigation graphs, and a batch of logging and filtering improvements across the platform.

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

01 **Filter enhancements in technical views** NEW 60

Adds a `revoked` filter option in technical views for Courses of Action and related entities, and adds a `CURRENT ENTITY` option in ID filters of 'in regards of (dynamic)' sub-filters, enabling self-referential dynamic filter rules.

– Names both filter options exactly, no usage walkthrough 7.260907.0

THINNER COVERAGE BELOW

02 **FIPS 140-3 compliant deployment images** NEW 52

OpenCTI now provides FIPS 140-3 compliant container images for regulated deployments, adopting the FIPS 140-3 validated `python-nodejs-fips` base image for the Docker deployment to run OpenCTI in FIPS 140-3 compliant mode.

– Names the base image but no deployment steps given 7.260907.0

03 **Inline entity creation in investigation graphs** NEW 50

Adds a Create Entity button to the investigation graph and the add-entity drawer, enabling analysts to create new entities and link them into the current investigation without leaving it.

When building an investigation graph and you realize a relevant entity is missing, use the new Create Entity button to add and link it without navigating away.

> 📍 In the investigation graph, open the Add entity drawer, then click Create Entity to define a new entity directly and have it added to the current investigation.

– Clear UI path and worked example, no deeper mechanism 7.260907.0

04 **Ingestion and connector activity logging** NEW 43

Adds ingestion logs for CSV feeds, surfacing feed-level activity for operators monitoring data ingestion, and adds activity logging for managed connector configuration updates.

— Names two logging surfaces but no access details 7.260907.0

05 **New workflow UI for entity types** NEW 18

Adds a new workflow UI applied across all entity types.

— Bare one-line mention with no detail on behaviour 7.260907.0

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ 3 RELEASES · 2026-09-07 BLOG ↗

RANK

Trend Vision One is Trend Micro's cloud-based XDR platform providing unified security visibility and threat detection across endpoints, networks, and cloud environments.

> Trend Vision One added an automated migration path from Worry-Free agents to TrendAI Vision One Endpoint Security, taught Security Playbooks to automatically remediate container image vulnerabilities, and introduced a new predefined cloud app category for Zero Trust Secure Access rules.

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

01 **Automatic Worry-Free agent migration** NEW 65

Managed Worry-Free agents assigned to a Standard Endpoint Protection Manager and enrolled in Endpoint Security Policies can now be automatically upgraded to TrendAI Vision One™ Endpoint Security via Endpoint Security → Endpoint Inventory, with an option to migrate default group and policy settings from Worry-Free so existing configurations are preserved during the upgrade.

Migrate Worry-Free agents to TrendAI Vision One Endpoint Security while preserving existing group and policy configurations.

> 📍 In the Trend Vision One console, go to Endpoint Security → Endpoint Inventory, then select Update from Worry-Free and choose to migrate default group and policy settings from Worry-Free.

— Names exact UI path and migration option, but no limits or timing given.

Automatically update Worry-Free agents …

02 **All untrusted cloud apps category in Zero Trust Secure Access** NEW 65

Zero Trust Secure Access adds a predefined All untrusted cloud apps custom cloud app category that auto-updates to include every cloud app with an untrusted status, usable alongside existing All trusted cloud apps and All poor reputation cloud apps categories in Internet Access rules and risk rules.

Block all untrusted cloud apps for a user group via an Internet Access rule without manually maintaining the app list.

> 📍 In the console, go to Zero Trust Secure Access → Secure Access Resources → Custom Cloud App Categories, confirm All untrusted cloud apps is listed, then

navigate to your Internet Access rules and add `All untrusted cloud apps` as a destination condition with a Block action.

— Names exact category values and where they're used, but limited mechanism detail.

Zero Trust Secure Access adds the All u…

03 **Automated container vulnerability response in Security Playbooks**

NEW 63

Security Playbooks, found under Workflow and Automation → Security Playbooks, gained container image vulnerability response actions that let a playbook automatically notify the security team, submit a ticket to a third-party ticketing system, quarantine the affected container, and apply remediation steps without manual intervention.

Build a playbook that automatically quarantines a vulnerable container and opens a ticket the moment Trend Vision One detects a container image vulnerability.

📍 In the console, go to Workflow and Automation → Security Playbooks, create a new playbook, set the trigger to a container image vulnerability detection event, then add actions in sequence: `Notify security team`, `Submit ticket to ticketing system`, `Isolate affected container`, and `Apply remediation steps`.

— Concrete action sequence and console path, but no trigger/detection detail.

Security Playbooks automate mitigation …

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗ RANK

Give feedback to the Intruder team so we can make more informed product decisions.
Powered by Canny.

> Intruder authenticated web application scans can now follow authentication and API flows across additional domains, including subdomains and third-party services.

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

01 Additional domains for authenticated web app scans NEW

80

Authenticated web app target configuration gains an Additional domains field that lets you include subdomains (e.g. `api.example.com`, `auth.example.com`) and external services such as identity providers, CDNs, and third-party APIs in scan scope. This enables cross-domain crawling and testing during authenticated scans so API backends and external auth providers are no longer blocked at the domain boundary.

When your app at `app.example.com` calls an API at `api.example.com` and authenticates via an external IdP, add those domains so the scanner can follow the full application flow.

> 📍 In the Intruder console, open your web app target › Authentication settings › locate the Additional domains field › add entries such as `api.example.com` and `auth.example.com` › save and re-run the authenticated scan.

— Names exact config field, example domains, and UI navigation steps.

Scan additional domains for authenticat…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Cohesive VNS3 Controller

Feedback ⓘ

1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Cohesive VNS3 Controller is a virtual network security appliance that orchestrates encrypted network connections and provides centralized management for hybrid cloud environments.

> VNS3 Controller's snapshot-20260907 release adds idempotent tunnel bounces, strict CIDR route validation, and licensing support for new connection types and firewall objects, while a docs update adds OpenLDAP-based user account management; the release also carries a breaking OpenVPN server upgrade to 2.6+.

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

01 **Optional plugin recovery on snapshot restore** NEW 90

Adds `recover_plugins` argument to the snapshot-restore API, defaulting to `true`, with a corresponding UI tick box (defaulting to `false`) to make plugin recovery from configuration snapshots optional.

— Names exact API argument, UI control, and default values snapshot-20260907

02 **Breaking OpenVPN server upgrade to 2.6+** BREAKING 75

Upgrading to VNS3 7.0.1 upgrades the OpenVPN server to version 2.6+; significant configuration differences between OpenVPN 2.5 and 2.6 may require manual modifications to existing OpenVPN clientpack hosts for proper reconnection — testing before production upgrade is recommended. Wireguard clientpacks are unaffected.

— Names exact versions and impact but no migration steps given snapshot-20260907

03 **Idempotent tunnel bounce for IPsec reset** IMPROVED 65

Tunnel bounce (IPsec reset) calls are now idempotent: if a bounce command for a tunnel is already queued, no additional bounce commands for that tunnel are enqueued until the existing one is processed.

— Describes mechanism clearly, no API/flag name given snapshot-20260907

04 **Licensing system tracks new connection types and firewall objects**

NEW 60

VNS3's licensing system now covers EdgeGRE, VXLAN, and Geneve connection types, plus firewall object counts spanning rules in main and sub-tables, sub-tables themselves, and firewall sets (`fwsets`).

— Names all licensed items but no command or endpoint to check them `snapshot-20260907`

THINNER COVERAGE BELOW

05 **Strict CIDR parsing for route entries** IMPROVED 55

Route entry now uses strict CIDR parsing: any route whose CIDR notation resolves to a different network than intended is rejected outright rather than silently accepted.

— Explains validation behaviour but not the error message or UI location

`snapshot-20260907`

06 **Status Dashboard as default view** IMPROVED 55

Sets the Status Dashboard as the default status view (replacing the Status Table) on systems with 2 vCPUs and 1 GB or more of memory.

— Names exact hardware condition and UI change `snapshot-20260907`

07 **UI edit panel for user-managed routes** IMPROVED 50

Adds a UI edit panel for user-managed route entries, replacing the previous delete-and-re-enter workflow.

— Names the change but not the exact UI path `snapshot-20260907`

08 **Wireguard background-process efficiency at scale** IMPROVED 50

Improves Wireguard background-process efficiency so CPU usage scales more efficiently beyond 1,000 clientpacks.

— Gives a concrete scale threshold but no benchmark numbers `snapshot-20260907`

09 **OpenLDAP integration for user accounts** NEW 40

VNS3 Controller adds OpenLDAP integration to manage VNS3 user accounts through a directory service.

— Names the integration but no config steps or fields `product docs`

10 **Fixed-size logging partition** IMPROVED 35

Adds a fixed-size logging partition so log overflow cannot impede Web UI or API operation.

— States purpose but no size or configuration detail `snapshot-20260907`

11 **HTTPS certificate upload before licensing** IMPROVED 35

Uploads of HTTPS certificates via the UI are now allowed before the controller is licensed.

— Clear behaviour change but minimal detail `snapshot-20260907`

└──▷ **BREAKING ON UPGRADE**

! Upgrading to VNS3 7.0.1 upgrades the OpenVPN server to version 2.6+; significant configuration differences between OpenVPN 2.5 and 2.6 may require manual modifications to existing OpenVPN clientpack hosts for proper reconnection — testing before production upgrade is recommended. Wireguard clientpacks are unaffected.

WAS THIS USEFUL?  

Cisco Duo Desktop

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07

NOTES ↗ RANK

Cisco Duo secures workforce access and defends against advanced identity threats with IAM, phishing-resistant MFA, SSO, passwordless, and device trust.

> Duo Desktop adds an optional telemetry consent flow with an admin override, smarter password-lockout prevention tied to Microsoft Entra ID join status, and several preferences and update-tracking improvements.

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

01 **Startup error/usage data collection consent** NEW 55

Duo Desktop now shows an optional error and usage data collection consent prompt at startup, which users can enable or disable at any time in the app's preferences. Administrators can set a system-wide preference to bypass this startup prompt for all users.

– Names both user and admin controls but no config key given snapshot-20260907

02 **Implicit password detection via Entra ID join status** NEW 50

Adds support for implicit password detection based on Microsoft Entra ID join status, preventing lockouts from health checks.

– Explains mechanism and benefit but no user-facing action snapshot-20260907

03 **Update status shown in preferences page** IMPROVED 45

The preferences page now shows whether the app is up to date and provides a link to download the latest version when it is not.

– Clear UI location but no exact path given snapshot-20260907

04 **Release track recorded in install directory** NEW 30

Duo Desktop now writes its release track to the application's install directory.

– Names surface but no path or use case snapshot-20260907

05 **Groundwork for Entra ID Trusted Endpoints and Secure Access** NEW 20

Internal changes made to support an upcoming feature leveraging Microsoft Entra ID for Trusted Endpoints and Secure Access.

– Preview of future feature with no current functionality snapshot-20260907

06 **New and updated translations** IMPROVED 5

Adds new translations covering new and updated screens throughout the application.

— Bare mention with no specifics `snapshot-20260907`

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Didit

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit's largest addition this window is per-country ID verification method configuration — letting operators independently enable document capture, non-document national-ID lookup, and digital ID wallets (MitID, BankID, itsme, and others) per country with fallback rules — alongside a new event-based transaction-monitoring capability, finer workflow routing, webhook assurance-tier fields, and QR/barcode decoding in document review.

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

01 Non-document national ID lookup verification method NEW

88

Non-document lookup checks a user's national ID number and personal details against a government register, credit-bureau, financial-services, utility, or residential-records source per country — no photograph or document upload required. Config keys `id_lookup.enabled`, `id_lookup.max_attempts`, `id_lookup.skip_liveness_and_face_match`, `id_lookup.on_partial_match`, `id_lookup.on_no_match`, and `id_lookup.on_provider_error` control availability, retry count (1–5, default 2, counting only attempts the source actually answered), and per-failure routing. In countries where the register returns a portrait (South Africa, Nigeria, Argentina, Panama), automatic selfie capture, passive liveness, and face match run so downstream Liveness and Face Match steps can be skipped. Session results gain a Registry comparison panel (compared fields, provided vs. source values, match verdicts, source name, timestamp, attempt count) and a See details view of the expected request/response for each lookup source.

Configure non-doc lookup attempts and fallback rules per country so users fall back to document capture only after exhausting lookup options.

📍 In the Didit console, go to ID Verification › Countries and allowed ID methods, select a country, then open Additional settings and set lookup attempts (1–5) and enable fallback to document capture.

Enable non-document national-ID lookup for a specific country so users can verify without uploading a document photo.

📍 In the console, open your workflow → ID Verification step → Countries tab, select the target country, and toggle Non-document lookup on. Set retry count (1–5,

default 2) and choose the fallback action for each failure mode (partial match, no match, no source response).

— Full mechanism, config keys, retry defaults, and result-panel fields named.

snapshot-20260907

02 **Digital ID wallet verification method** NEW 79

Adds digital ID wallets as an ID Verification method — MitID, BankID, itsme, Vipps, iDIN, UAE PASS, gov.br, and EUDI Wallet — configurable per country via `wallet.enabled`, `wallet.providers`, and `wallet.on_failure`. Wallets started as Coming soon on the roadmap; production wallets are now enabled live while sandbox wallets can be tested with simulated results and no charges. Session results gain a Credential details /'Credential' card showing the wallet used, issuing authority, credential type, assurance level, signature verification outcome, and shared attributes.

Set per-country ID verification methods and fallbacks, including digital ID wallets like BankID or EUDI Wallet, for a specific country.

📍 In the console, go to ID Verification Methods › Countries tab, select a country, enable digital ID wallet methods (e.g. BankID, EUDI Wallet), and configure fallback to document capture.

— Named providers and config keys, but no runnable endpoint given. snapshot-20260907

03 **Per-country ID verification method controls** NEW 70

Adds a Countries tab on the ID Verification node in the workflow console with a Global settings bulk switch that enables a method across all supported countries at once, letting operators independently enable or disable document capture, non-document lookup, and digital ID wallets per country via the console or API. Document capture is governed by `document.enabled` and `documents_allowed`; an Additional settings panel adds document subtypes and fallback rules. A `fallback_from` field is added to session records, populated when a country's fallback switch is set to decline and the session was declined.

— Config keys and UI tab named, but no worked example attached. snapshot-20260907

THINNER COVERAGE BELOW

04 **Event-based transaction monitoring** NEW 57

Adds SDK event submission for device-side capture of non-monetary events using a scoped monitoring credential, with event ingestion for types including login, gameplay, bonus, payment, dispute, security, and responsible gambling. Adds rule types on events supporting velocity, ratio, ordered-sequence, and cross-subject correlation.

— Names event and rule types but no concrete API or config given. product docs

05 Verification method and assurance tier in status webhook NEW

52

Adds `status.updated` webhook fields carrying the verification method used (Documentary , Data match , or Cryptographic) and the corresponding assurance tier for every ID Verification session.

— Names the webhook and enum values but no usage example. snapshot-20260907

06 Workflow routing by document attributes and risk signals

IMPROVED

50

Workflows can now route users by `document subtype` , date of issue, age, months since issue, risk score, webhook JSON response paths, and custom rules, in addition to existing country, document type, and verification status routing.

Route users in a workflow by document subtype or risk score to apply different verification paths based on finer identity signals.

> 📍 In the Didit console workflow builder, add a routing rule and select document subtype , risk score , or months since issue as the condition field to branch the user journey.

— Named routing fields but no mechanism beyond the list. product docs

07 QR code and barcode decoding in document review NEW

50

Session view now highlights and decodes QR codes and barcodes found on scanned documents: Document AI marks the code's location on the image or PDF preview and shows the decoded payload; codes that fail to decode are shown as detected-only with no payload.

When reviewing a document session, inspect a detected QR code or barcode — the location overlay and decoded payload appear separately from any document-manipulation warning highlights.

> 📍 In the Didit Console, open a verification session › navigate to the Document AI results › locate the QR/barcode overlay on the image or PDF preview to view the decoded payload, or note detected only if decoding failed.

— Describes behaviour and UI path but no config surface. product docs

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Ubiq Security

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗

RANK

Learn how Ubiq discovers sensitive data, maps identity-to-data access, and enforces protection policies across applications, databases, APIs, and AI.

> Ubiq shipped native on-database column-level encryption for Postgres and Greenplum, extended Partial Encryption to six client libraries plus the Dashboard, and rounded out the release with search-ready encryption, billing metadata support, and expanded security history logging.

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

01 Partial Encryption across client libraries and Dashboard NEW

82

Adds Partial Encryption support for C/C++, Go, Python, Java, C#.NET, and NodeJS Client Libraries (version 2.0.0+) to mask or partially encrypt Structured Data, plus a Partial Encryption option for Structured Datasets in the Dashboard. Also adds a Failed Key Retrieval Attempt error to the security history log to surface failed Partial Encryption attempts from unsupported Client Library versions (< 2.2.0).

Enable Partial Encryption for a Structured Dataset so sensitive columns are masked for unauthorized readers while remaining searchable.

> 📍 In the Dashboard, navigate to Datasets › select your Structured Dataset › Edit › enable the Partial Encryption option, then save.

— Names every SDK, versions, and includes Dashboard navigation steps snapshot-20260907

THINNER COVERAGE BELOW

02 On-database encryption for Postgres and Greenplum NEW

43

Adds native integration for Postgres and Greenplum enabling on-database column-level encrypt/decrypt directly from SQL.

— Names the databases and SQL-level scope but no exact syntax snapshot-20260907

03 Security history log gains rate limit and account change events

IMPROVED

40

Adds rate limit event logging to the Security History log, visible to the affected user, and adds security history logging for billing contact and account name change events.

— Named log events, no navigation or example given snapshot-20260907

04 **EncryptforSearch support in Go Client Library** NEW

30

Adds `EncryptforSearch` support to the Go Client Library.

— Named addition only, no usage or mechanism given snapshot-20260907

05 **Custom metadata in billing records via Python wrapper** NEW

30

Adds a wrapper capability for the Python Client Library enabling customer-defined metadata in billing records.

— No code sample or field names given snapshot-20260907

WAS THIS USEFUL? 👍 👎

Immuta

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Use Immuta to provision governed access to data.

> Immuta introduces public preview intent-based access control for Databricks Unity Catalog and extends subscription policies to metric views, while tightening lockout behavior for policies referencing missing columns.

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

01 Intent-based access control for Unity Catalog NEW

67

Introduces intents (public preview) for Databricks Unity Catalog: policy authors configure groups authorized to use an intent and attributes that replace a user's birthright attributes while the intent is active, enabling use-case-driven, auditable access to sensitive data such as PII.

– Explains mechanism and purpose but no UI/API steps given snapshot-20260907

THINNER COVERAGE BELOW

02 Lockout policy applies to references to missing columns

55

BREAKING

Immuta will apply a lockout policy to data sources governed by a policy that references a nonexistent column or column tag; previously no lockout applied in these cases. Policies referencing missing columns or column tags will now trigger a lockout on rollout.

– Clear before/after behavior change but no migration steps snapshot-20260907

03 Subscription policies for Unity Catalog metric views NEW

43

Adds subscription policy support for Databricks Unity Catalog metric views, enabling access provisioning and governance for that object type through existing Immuta policies.

– States capability but lacks configuration detail snapshot-20260907

└──▷ **BREAKING ON UPGRADE**

! Immuta will apply a lockout policy to data sources governed by a policy that references a nonexistent column or column tag; previously no lockout applied in these cases. Policies referencing missing columns or column tags will now trigger a lockout on rollout.

WAS THIS USEFUL?  

Securiti

Feedback ⓘ 1 RELEASE · 2026-09-07 BLOG ↗ RANK

Securiti DataAI Command Platform™: Empowering safe use of data and AI. Achieve unified intelligence and controls across hybrid multicloud environments, ensure security, privacy, governance, and compliance.

> Securiti launched Consent Agent, an agentic Build-Assess-Remediate system that automates cookie consent configuration, testing, monitoring, and audit evidence across 80+ jurisdictions.

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

01 Automated compliance testing across 20+ checks NEW

75

Consent Agent runs automated compliance testing across 20+ checks covering banner correctness, pre-consent tracking, TC string validity, GPC signal handling, dark pattern detection, accessibility, and privacy policy completeness. Each check is mapped to a specific legal requirement and returns pass/fail/NA results with jurisdiction-aware risk and confidence scoring.

— Names every check type and scoring mechanism, no runnable interface given

Cookie Consent Meets Agentic Automation…

02 Consent Agent Build-Assess-Remediate workflow NEW

65

Consent Agent is a three-stage agentic workflow — Build, Assess, Remediate — that automatically configures jurisdiction-correct consent banners, runs 20+ compliance tests, and applies or escalates fixes across all domains.

— Explains the three stages but no exact commands or config surfaces

Cookie Consent Meets Agentic Automation…

03 Centralized compliance dashboard NEW

65

A centralized compliance dashboard offers pre-built widgets surfacing pass/fail status across domains, tests, and frameworks (IAB EU/US, GPC, Google Consent Mode, GDPR, CCPA), an aggregate compliance score, trend graphs, and a prioritized findings list with per-risk assignment, acceptance, exclusion, and project-tracking.

— Names frameworks and widgets but only a UI path, no exact navigation given

Cookie Consent Meets Agentic Automation…

THINNER COVERAGE BELOW

04 Continuous domain rescanning for consent drift NEW

50

Domains are rescanned on a schedule so configuration drift — new pixels, misfiring tags, pre-consent script loads — is caught and flagged against a specific test and point in time rather than discovered at the next audit.

— Describes mechanism but no schedule interval or config detail

Cookie Consent Meets Agentic Automation…

05 **Consent Regulatory Database** NEW 45

A Consent Regulatory Database backs Consent Agent with out-of-the-box cookie consent templates for 80+ jurisdictions, consent requirements, detailed cookie guidelines, and case law.

— Names scope (80+ jurisdictions) but no interface or usage detail

Cookie Consent Meets Agentic Automation…

06 **Audit logs for remediation actions** NEW 40

Consent Agent generates before-and-after audit logs for every automated and manual remediation action, producing timestamped, audit-ready evidence for regulators and internal audits.

— States the output but no format, retention, or access detail

Cookie Consent Meets Agentic Automation…

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Promptfoo

Feedback ⓘ 1 RELEASE · SEEN 2026-09-07 NOTES ↗ RANK

Promptfoo tests, evaluates, and red-teams LLM applications to detect vulnerabilities and compare model outputs.

> Promptfoo expanded its provider ecosystem with new integrations for Docker-isolated model execution, MCP servers, multimodal image models, and additional reasoning models, alongside new client-certificate support for HTTP-based providers.

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

01 **JKS/PFX client certificates for HTTP providers** NEW 62

HTTP provider now supports JKS and PFX client certificates, enabling mTLS connections to secured endpoints.

– Names cert formats and mTLS but no config key shown snapshot-20260907

THINNER COVERAGE BELOW

02 **New model provider integrations** NEW 55

Adds several new providers and models: Docker Model Runner for running models in isolated Docker containers for better security and reproducibility; Google Imagen for image generation in multimodal testing scenarios; AIMLAPI for accessing various AI models through a unified interface; Grok-4 via xAI for advanced reasoning evaluations; `o3-deep-research` and `o4-mini-deep-research` OpenAI Deep Research models for complex problem-solving evaluations; LiteLLM Embeddings support for similarity testing and semantic search evaluations; and Google Vision for image understanding in multimodal evaluations.

– Names each provider but gives no config or usage steps snapshot-20260907

03 **MCP provider for connecting to MCP servers** NEW 50

Adds an MCP (Model Context Protocol) provider to connect to MCP servers, including a red-team agent testing example with tool call results.

– Describes capability and example but no setup steps snapshot-20260907

04 **System prompt support in Azure provider** IMPROVED 30

Enhances the Azure provider with system prompt support for better control over model behavior.

– Bare mention of enhancement, no mechanism or example snapshot-20260907

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ 1 RELEASE · 2026-09-07

NOTES ↗ RANK

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

> SkillSpector v2.11.1 makes its scan workflow deadline configurable and raises the default from 60 to 600 seconds, reducing timeouts on larger recursive scans.

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

01 Configurable workflow deadline via environment variable

IMPROVED 86

Adds the `SKILLSPECTOR_MAX_WORKFLOW_SECONDS` environment variable to set a custom positive finite aggregate workflow deadline for the end-to-end workflow and transitive traversal; invalid, zero, negative, infinite, and NaN values fall back to the new 600-second default (raised from the previous 60 seconds). The configured deadline is applied consistently across direct CLI, recursive, transitive, and multi-skill analysis paths, preventing larger valid scans from timing out.

Override the default 600-second deadline when scanning a large recursive skill tree that needs more time.

```
$ SKILLSPECTOR_MAX_WORKFLOW_SECONDS=1800 skillspector scan ./my-large-skill/
```

Set a persistent custom deadline in a Docker env-file so all container-based scans share the same budget.

```
$ SKILLSPECTOR_PROVIDER=anthropic
ANTHROPIC_API_KEY=sk-ant-...
SKILLSPECTOR_MAX_WORKFLOW_SECONDS=1200
```

— Names exact env var, old/new defaults, and fallback behaviour with runnable examples.

v2.11.1

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ 1 RELEASE · 2026-09-07 NOTES ↗

RANK

Docker Sandboxes isolates applications in lightweight containers, providing secure, reproducible environments for development and deployment.

> Docker Sandboxes' v0.42.0 window adds cloud-hosted sandboxes, declarative `sbxenv.yaml` environment files with lifecycle hooks and parameterized args, native Devin agent support, a persistent shared skills store, and expanded kit customization (arguments, OCI/Git kit references, mixins), alongside registry mirror routing, mountless sandboxes with named reattachment, MCP Gateway SSRF/scope controls, and a machine-wide Windows installer — plus breaking changes to default port binding and `.sbxenv.yaml` variable interpolation.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 **Declarative sbxenv.yaml environment files** NEW 100

Introduces `sbxenv.yaml` files that declaratively define a sandbox — workspaces, kits, ports, secrets, bindings, MCP registrations and sandbox options — managed via new `sbx env plan [PATH...]` (preview changes without applying), `sbx env create --auto-approve` (apply unattended), `sbx env run --env-arg NAME=VALUE` / `--env-args-file <file>`, `--skip-host-commands`, and `--force` (approve destroy of an in-use sandbox). A top-level `args` block supports `default`, `required`, `enum`, `pattern`, `description` and `channel` fields referenced as `${{ env.args.NAME }}`; a `lifecycle` block defines `initialize`, `postCreate` and `preRemove` phases with `cmd`, `label`, `workdir` and `timeout`, exposing `SBX_LIFECYCLE_PHASE`, `SBX_ENV_FILE`, `SBX_ENV_FILES`, `SBX_ENV_DIR`, `SBX_SANDBOX_NAME`, `SBX_AGENT` and `SBX_WORKSPACE` to host commands. A user-level `~/.sbxenv.yaml` merges beneath the project file as shared defaults; `sbxenv.yaml` is now read in preference to `.sbxenv.yaml` when both exist. Sandbox options include `shareSkills: false`, `display`, `gpu` and `usb`; `env.rememberHostCommands` records approval of lifecycle/credential host commands until the command text changes, and `sbx env create` now shares imported skills by default.

Preview exactly what an environment file will change on the host before committing — useful before a first apply or after editing `sbxenv.yaml`.

```
$ sbx env plan ./web-app-env/sbxenv.yaml
```

Supply per-run argument values (e.g. target a beta release channel and a specific API endpoint) without editing the shared environment file.

```
$ sbx env run --env-arg channel=beta --env-arg
endpoint=https://api.example.com
```

Parameterize a sandbox environment so the same file can target different release channels and API endpoints, with validation enforced at apply time.

```yaml
args:
  channel:
    default: stable
    description: Release channel
    enum:
      - stable
      - beta
  endpoint:
    required: true
    description: API endpoint
    pattern: "https://.*"

env:
  RELEASE_CHANNEL: ${{ env.args.channel }}
  API_ENDPOINT: ${{ env.args.endpoint }}

lifecycle:
  initialize:
    - label: Prepare workspace
      cmd: test -d web-app || git clone
https://github.com/example/web-app
      timeout: 5m
```

— Extensive named blocks, flags and env vars; mechanism thoroughly explained. v0.42.0

02 **Kit arguments and remote kit references** NEW 95

Adds `--kit-arg name=value` and `--kit-arg <kit-name>.name=value` (prefixed to target a specific kit, taking precedence over unprefixed values) to `sbx run` / `sbx create`, supplying typed inputs to schema v2 kits declaring an `args:` block, interpolated via `${{ kit.args.<name> }}`; adds `--kit-args-file <path>` to load reusable `name=value` files (blank lines/`#` comments ignored, later files and explicit `--kit-arg` override). Supports launching kits from OCI registries (`sbx run docker.io/sbx/droid-kit:latest`) and Git repos (`git+https://github.com/docker/sbx-kits-contrib.git#ref=v0.1.0&dir=amp`), plus local directory/ZIP paths starting with `./` or `../`. From v0.42.0, `sbx run <sandbox-kit-ref> PATH...` and `sbx`

`create <sandbox-kit-ref> PATH...` take the kit reference positionally, deprecating the old `sbx run <name> --kit <ref>` form — `--kit` (repeatable) is now used for attaching mixins. The same `--kit-arg` / `--kit-args-file` flags apply to `sbx kit validate` and `sbx kit inspect`.

Supply per-kit arguments from a file when creating a sandbox — useful for reproducible CI pipelines where secrets are kept out of shell history.

```
$ sbx create ./my-agent/ . --kit ./my-mixin/ --kit-args-file
./kit.args --kit-arg my-mixin.channel=stable
```

Launch a versioned sandbox kit directly from an OCI registry, targeting a specific tag.

```
$ sbx run docker.io/sbx/droid-kit:latest
```

— Names every flag, precedence rule and command; mechanism fully explained. v0.42.0

---

03 **Registry mirror for Docker Hub pulls** NEW 88

Adds `platform.images.registryMirror` (set via `sbx settings set platform.images.registryMirror <host>`, removed via `sbx settings unset platform.images.registryMirror` or an empty value) to route Docker Hub pulls for sandbox templates and OCI kits through an organization's registry mirror, with optional port and path-prefix support; when the value is a non-loopback host with no path prefix, the mirror is also propagated to Docker Engine inside the sandbox so in-sandbox `docker pull` uses it. The same mirror can be declared directly in a `registry mirror` block in `sbxenv.yaml`.

Route all sandbox template, OCI kit, and in-sandbox Docker Hub pulls through a corporate registry mirror.

```
$ sbx settings set platform.images.registryMirror
registry.example.com
```

Mirror sandbox template and OCI kit pulls through a registry that stores Docker Hub content under a repository path prefix (in-sandbox Docker Engine will not use this mirror).

```
$ sbx settings set platform.images.registryMirror
registry.example.com/docker-remote
```

Remove the registry mirror configuration so sandbox templates and OCI kits pull directly from Docker Hub.

```
$ sbx settings unset platform.images.registryMirror
```

— Exact setting keys and commands given; mirror propagation mechanism explained. v0.42.0

04 **MCP Gateway SSRF checks and auth scope controls** NEW

88

Adds `--skip-ssrf-check` to `sbx mcp add` to suppress the SSRF warning/check when registering a trusted internal MCP server resolving to a private, loopback, link-local or cloud-metadata address, and now allows `--url` on such addresses with a warning instead of outright rejection. Adds `--skip-auth` to `sbx mcp add` (the older `--skip_auth` still works). Adds `--scope` to `sbx mcp auth` to override recorded default scopes for a single authorization (resolution order: `--scope` argument, `sbx mcp add --scope` default, protected-resource required scopes, then none) and `--no-scope` to suppress both recorded and required scopes (cannot combine with `--scope`). Adds `--json` to `sbx mcp auth status` reporting granted, default, supported, duplicate and unexpected scopes.

Register a trusted internal MCP server that resolves to a private IP without triggering the SSRF warning, so internal tooling can be wired up without noise.

```
$ sbx mcp add myserver --url https://internal.corp/mcp --skip-ssrf-check
```

Authorize against an OAuth-backed MCP server with a narrower scope than the recorded defaults — useful when a workflow only needs read access.

```
$ sbx mcp auth serverx --scope read
```

Pull machine-readable auth-status JSON for the Notion MCP server to feed into a CI check or audit script.

```
$ sbx mcp auth status notion --json
```

— All flags and scope-resolution order named with runnable examples. v0.42.0

05 **Native Devin agent support** NEW

80

Adds `sbx run devin` to launch a Devin agent in a sandbox (workspace path optional, defaults to current directory), built on the `docker/sandbox-templates:devin-docker` base image which can be customized to pre-install tools. Reads `AGENTS.md` from the workspace for agent-specific instructions and starts Devin with `--permission-mode dangerous --respect-workspace-trust=false` by default. Devin credentials are captured after first interactive sign-in and reused automatically for future Devin sandboxes through a proxy.

Run Devin against a specific project directory in an isolated sandbox — useful for CI or one-off agent tasks on a repo you don't want touching your host environment.

```
$ sbx run devin ~/my-project
```

Run Devin from within the project directory itself, letting the workspace default to the current directory.

```
$ cd ~/my-project && sbx run devin
```

— Command, base image, defaults and credential flow all specified. v0.42.0

06

Mountless sandboxes, named reattachment and file transfer

NEW

78

Adds mountless sandbox creation via `sbx create --name <name> <agent>` (no host directory bind mount), reattachable later with `sbx run --name <name>` from any working directory without re-specifying the agent. Adds `sbx cp` to copy files between host and sandbox by absolute sandbox path, e.g. `sbx cp ./config.json my-sandbox:/home/agent/workspace/` and `sbx cp my-sandbox:/home/agent/workspace/output.log ./`. Appending `:ro` to a workspace path mounts an extra workspace read-only inside the sandbox.

Create a mountless scratch sandbox and attach to it later — useful when you want an isolated environment with no host files exposed.

```
$ sbx create --name scratch claude
sbx run --name scratch
```

Pull generated output or logs out of a running sandbox to the host without needing a mounted workspace.

```
$ sbx cp my-sandbox:/home/agent/workspace/output.log ./
sbx cp ./config.json my-sandbox:/home/agent/workspace/
```

— Commands and exact syntax given for each named capability. v0.42.0

07 **Agent skills store and management commands** NEW 75

Adds `sbx skills add` to install all skills from a Git repo (`owner/repository` shorthand or HTTPS/SSH URL), with `--skill` to install specific named skills (comma-separated, e.g. `--skill frontend-design --skill pdf`); `sbx skills ls` lists installed skills; `sbx skills update [names...]` refreshes all or named skills; `sbx skills rm` removes one or more. Skills live in a persistent shared store that survives sandbox deletion and is automatically shared with new sandboxes running a supported agent.

Install specific skills from a GitHub repository into the shared store so every new sandbox picks them up automatically.

```
$ sbx skills add anthropics/skills --skill frontend-design --
skill pdf
```

Refresh only the skills you actively use after upstream updates, without touching the rest of the store.

```
$ sbx skills update frontend-design pdf
```

Audit what is currently installed, then remove a skill no longer needed by running agents.

```
$ sbx skills ls && sbx skills rm frontend-design
```

— All subcommands and flags named; store persistence briefly explained. v0.42.0

08 **Windows MSI installers for machine-wide and per-user install**

NEW 70

Adds `DockerSandboxesMachine.msi` for silent, machine-wide Windows installation via `msiexec.exe /i DockerSandboxesMachine.msi /quiet` , installing to `%ProgramFiles%\DockerSandboxes` and adding its `bin` directory to the system `PATH` with no per-user install required. Per-user installation via WinGet

(`DockerSandboxes.msi`) places files in `%LOCALAPPDATA%\DockerSandboxes` and updates the user `PATH` without admin privileges.

Silently deploy Docker Sandboxes to all users on a managed Windows fleet from an elevated PowerShell prompt.

```
$ msiexec.exe /i DockerSandboxesMachine.msi /quiet
```

— Exact installer names, paths and commands given for both modes. v0.42.0

09 **JSON output for read-only commands and server state** NEW

70

Adds `--json` output to read-only commands `secret ls` , `version` , `mcp ls` , `skills ls` , `policy inspect` and the `kit` verification commands. `sbx version --json` now reports a `server.state` field (`running` or `unavailable`) so scripts can check backend reachability without parsing error text.

Check in CI whether the sbx daemon is reachable without parsing human-readable error text.

```
$ sbx version --json | jq -r '.server.state'
```

— Lists affected commands and new field, thin on mechanism. v0.42.0

10 **Published ports default to tcp4 only** BREAKING

70

`sbx ports --publish` and kit-declared ports now default to `tcp4` instead of dual-stack `tcp` , so a published port no longer listens on `::1` ; use `--publish 8080:3000/tcp` to restore dual-stack behaviour.

— Exact before/after behaviour and restoring flag given. v0.42.0

11 **Configurable sandbox volume size** NEW

65

Sandbox Docker volume size is configurable at creation time via the `DOCKER_SANDBOXES_DOCKER_SIZE` environment variable (minimum 512 MiB); v0.42.0 changes the default from 50 GB to 10 GB (sparse file) to reduce host disk usage. Existing volumes are not resized.

— Names env var and exact default change but limited scope. v0.42.0

12 **sbxenv.yaml interpolation and workspace defaults changed**

65

BREAKING

`${VAR}` interpolation in `.sbxenv.yaml` is no longer expanded — use the new `args:` block with `${{ env.args.NAME }}` syntax instead. An environment file that declares no `workspace:` now creates a sandbox with no workspace bind mount; write `workspace: .` to mount the project directory as before.

— Two named breaking changes with exact syntax and workaround. v0.42.0

---

13 **Cloud-hosted sandboxes via sbx --cloud** NEW 60

Adds `sbx --cloud` to run AI agents on Docker-managed cloud infrastructure, with `sbx --cloud run`, `sbx --cloud cp`, `sbx --cloud secret` and `sbx --cloud policy` covering file transfer, credentials and outbound network policy, plus `sbx move` to copy a sandbox filesystem between local and cloud environments.

Run a cloud-hosted Claude agent on Docker-managed infrastructure and give it a name for the project session.

```
$ sbx --cloud run claude --name cloud-project
```

— Names subcommands but omits internal mechanism of cloud infrastructure. v0.42.0

THINNER COVERAGE BELOW

14 **MCP policy governance and audit improvements** IMPROVED 30

A declined `@requireApproval` prompt on a local MCP server now generates its own audit record with policy attribution and a context digest, and governance resolution issues are now surfaced in `sbx policy` output and the dashboard.

— Describes outcome only; no named endpoint or config surface. v0.42.0

---

15 **Automatic diagnostics bundle upload on daemon error** NEW 15

Docker Sandboxes can now upload a diagnostics bundle automatically when the daemon hits an error, after opt-in.

— One-line description; no mechanism, trigger detail or format given. v0.42.0

---

16 **Clipboard sync from sandbox to host** NEW 15

Clipboard commands inside local sandboxes can now copy text to the host clipboard.

— Single sentence; no command or configuration surface named. v0.42.0

└──▷ **BREAKING ON UPGRADE**

! The `sbx run <sandbox-kit-name> --kit <sandbox-kit-ref>` form is deprecated as of version 0.42.0; use `sbx run <sandbox-kit-ref> PATH...` instead.

! `sbx ports --publish` and kit-declared ports now default to `tcp4` instead of dual-stack `tcp`, so a published port no longer listens on `::1`; use `--publish 8080:3000/tcp` to restore dual-stack behaviour.

! `${VAR}` interpolation in `.sbxenv.yaml` is no longer expanded; use the new `args:` block with `${{ env.args.NAME }}` syntax instead.

! An environment file that declares no `workspace:` now creates a sandbox with no workspace bind mount; write `workspace: .` to mount the project directory as before.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 7, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe** Feedback on this issue