THE CYBER TOOLCHAIN

NO. 076

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 16, 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 199 changes

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 40 of 42 tools

▾ // **THE BRIEF** HIDE

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

PROTECT

Contain an overdue KEV host without waiting for the patch window

vim ? help

Qualys TruRisk Eliminate isolates hosts with overdue CISA KEV vulnerabilities without an EDR agent in the path, and scores patch reliability so the usual "will this break production" stall has evidence behind it; ring-based rollout stages the blast radius. Veracode's Pipeline Scan now takes artifacts over 200 MB, so the monolith that silently skipped scanning gets scanned.

Qualys · Veracode

PROTECT

Give an AI agent its own identity instead of a human's credential

JumpCloud Agent Identities (Public Preview) issues OAuth credentials and group-based SSO access to non-human agents and authorizes them against MCP servers, so an agent's access can be scoped and revoked on its own rather than riding a service account nobody owns. Aembit adds CrowdStrike AIDR as inline allow/block/transform on MCP tool traffic, tied to the calling identity.

JumpCloud · Aembit

PROTECT

Let a coding agent run shell commands without trusting it with the machine

VS Code now runs agent terminal commands under OS-level file system and network isolation, which is the difference between reviewing every proposed command and letting the agent work. Docker Sandboxes adds tri-state access control on the shared skills store and namespaces MCP gateway secrets, so one sandbox's credentials stop leaking into another's.

code-oss · Docker Sandboxes

OFFENSIVE SECURITY

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-09-16

PRESS ↗ RANK

Horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

> NodeZero now sends validated exposure findings directly into CrowdStrike Falcon Next-Gen SIEM, letting security teams correlate them with other telemetry during investigations.

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

01 **CrowdStrike Falcon Next-Gen SIEM integration** NEW 67

NodeZero validated findings and exposure data now flow into CrowdStrike Falcon Next-Gen SIEM, available via the CrowdStrike Marketplace listing at `marketplace.crowdstrike.com/listings/horizon3-ai-nodezero/`. This lets teams correlate validated exposure data with endpoint, identity, cloud, and other security telemetry during SIEM investigations, incorporating it into existing security operations workflows.

When setting up the NodeZero integration to send validated findings into Falcon Next-Gen SIEM for your organization.

> 📍 In the CrowdStrike Marketplace, navigate to the Horizon3 AI NodeZero listing at `marketplace.crowdstrike.com` /listings/horizon3-ai-nodezero/, then follow the installation steps to enable the integration and begin ingesting NodeZero findings into Falcon Next-Gen SIEM.

— Names exact marketplace URL and integration but no config/API detail

launch-20260916-ed32f58f

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

1 RELEASE · 2026-09-16 | CHANGELOG ↗ | RANK

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

> GitGuardian's detection engine added new analyzers for Akamai and Docker Hub, expanded n8n host detection, reduced false positives for generic high-entropy secrets, and updated its GitLab, Figma, and OpenAI analyzers to track upstream API changes.

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

01 **Analyzer updates for GitLab, Figma and OpenAI** IMPROVED

40

GitGuardian updated its `GitLab` secret analyzer to adapt to changes in GitLab API behavior, its `Figma Personal Access Token` analyzer to reflect changes in the Figma API, and its `OpenAI API Key` analyzer's scope checker endpoints following OpenAI API updates.

— Names the three analyzers updated but no user-facing action.

changelog-20260916-84cb4e9d

02 **New Akamai, Docker Hub and n8n detection** NEW

25

The detection engine added analyzers for `Akamai` and `Docker Hub`, expanded host detection for `n8n`, and reduced false positives for generic high-entropy secrets.

— Named only in a summary line, no mechanism or detail given. changelog-20260916-84cb4e9d

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗ RANK

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical v0.165.11 introduces a credential-free agent vault preview, adds Snowflake support to PAM, and rounds out the release with certificate, gateway, billing, and secrets-UI improvements.

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

01 **Copy secrets workflow in UI** NEW 50

Implements a copy secrets workflow in the secrets UI, enabling bulk secret copying across environments or projects.

– Names the UI location and bulk-copy behaviour clearly v0.165.11

02 **Credential-free agent vault (preview)** NEW 45

Introduces agent vault, a mode allowing Infisical agents to run without credentials; the capability is currently in preview.

– Names the feature and preview status but no config or CLI detail v0.165.11

03 **Optional relay for gateway deployments** IMPROVED 45

Makes the relay optional for gateways, reducing infrastructure requirements for gateway deployments.

– Explains the change but not the mechanism or config toggle v0.165.11

04 **Snowflake account support in PAM** NEW 40

The Privileged Access Management module now supports Snowflake accounts as a new account type.

– Names integration and module but no setup steps given v0.165.11

05 **Custom X.509 extensions for certificates** NEW 40

Certificate policies and profiles now support custom X.509 extensions.

– Names the extension type but no field or API detail v0.165.11

06 **Self-serve plan upgrades in UI** NEW 40

Adds self-serve plan upgrade capability directly from the product UI.

– Clear UI starting point but no further mechanism described v0.165.11

07 **Time-unit secret scanning environment variables** IMPROVED

Secret scanning environment variables now support time-unit values, enabling human-readable duration configuration instead of raw numbers.

— No specific env var names given despite naming the surface type `v0.165.11`

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ 2 RELEASES · 2026-09-15 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 added enterprise-level enforcement for Advanced Security configurations and introduced a Copilot-powered suggestion feature for custom repository property values.

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

01 Copilot suggests custom property definition values NEW

70

GitHub Copilot now suggests relevant allowed values when admins create custom property definitions (single-select or multi-select) for repositories, speeding up governance taxonomy setup. A new Repository custom property suggestions Copilot policy at the enterprise and organization level controls whether this suggestion behavior is enabled, available in public preview for Business and Enterprise plans.

Enable or restrict the new suggestion feature org-wide before rolling out custom property definitions to your repository fleet.

📍 In your organization (or enterprise) settings, go to Copilot › Policies and set Repository custom property suggestions to Enabled or Disabled.

— Names the policy and gives exact settings path to enable it.

GitHub Copilot suggests custom properti…

THINNER COVERAGE BELOW

02 Enforce Advanced Security configurations at enterprise level

NEW

55

A new Enforce for repository and organization owners option in the Enforcement dropdown of security configurations blocks organization and repository administrators from overriding enterprise-level GitHub Advanced Security settings.

— Names the exact dropdown option but no rollout mechanics.

Enforce GitHub Advanced Security config…

WAS THIS USEFUL? 👍 👎

Phoenix Security

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗

RANK

Phoenix Security correlates scanner, cloud and runtime findings with exploitability and ownership, and uses AI agents to prioritise and open remediation pull requests.

> Phoenix Security adds native scanner integration with Rapid7 InsightVM's SaaS platform for vulnerability ingestion.

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

01 **Rapid7 InsightVM scanner integration** NEW 88

Adds Rapid7 InsightVM (SaaS/cloud) as a new scanner integration, configurable via `Integrations > Scanners > Add Scanner Integration`. Select `Rapid7 Platform`, then supply a Server URL in the form `https://{region}.api.insight.rapid7.com/` (region codes: us, us2, us3, eu, ca, au, ap, aps2, me1), an API Username, and Password for a dedicated API user with MFA disabled, then click `Create Scanner` to pull vulnerability findings from monitored Sites into Phoenix environments.

Connect Phoenix Security to a Rapid7 InsightVM SaaS account so vulnerability findings from monitored Sites flow into Phoenix environments.

📍 1. In the sidebar, go to Integrations › Scanners. 2. Click `Add Scanner Integration`. 3. Select `Rapid7 Platform` from the integration list and enter a name. 4. Enter your Server URL (e.g. `https://us.api.insight.rapid7.com/` — replace `us` with your region code: us2, us3, eu, ca, au, ap, aps2, or me1). 5. Enter the API Username and Password for a dedicated API user (MFA disabled). 6. Click `Create Scanner`.

— Exact UI path, URL format and credential fields given

Rapid7 InsightVM Platform Integration (…

WAS THIS USEFUL? 👍 👎

Veracode

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

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities in code.

> Veracode's Static Analysis update adds a sweeping set of new language, framework, and platform versions (Dart, Flutter, PHP, .NET, Kotlin, JDK, Go, and more), an enhanced Python scanner with deeper taint tracking, Pipeline Scan support for artifacts over 200 MB, and new `veracode.json` exclusion controls for Java jars and classes.

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

01 **Enhanced Python scanner and taint analysis** IMPROVED 75

Adds an enhanced Python scanner with improved taint tracking, cross-file analysis, and pattern recognition covering PySpark, Flask, Django, and FastAPI. Also adds `send_from_directory` and `send_file` as taint sinks for CWE-73 detection, `httplib` support for detecting CWE-295 and CWE-297 SSL/TLS flaws, `os.path` propagation in taint analysis, and library support for `google-cloud-bigquery`, Psycopg 3, PyYAML, AWS Glue, and PySpark/DynamoDB/FastAPI.

— Names specific sinks, CWEs, and libraries with mechanism

Static Analysis updates (2026-09-03)

02 **.NET Semantic Kernel support and taint sinks** NEW 75

Adds .NET Semantic Kernel (SK) `InvocationFilter` and `Filter` support, SK `KernelContent` and Taint Injection support, and SK detection of hardcoded credentials (CWE-798) and configuration control issues (CWE-15). Also adds taint sink coverage for `System.Web.UI.Page.Title` and `System.Web.UI.HtmlControls.HtmlHead.Title` for CWE-80 detection in .NET.

— Names exact SK APIs, CWEs, and taint sink fields Static Analysis updates (2026-09-03)

03 **Gen IR `--capture` diagnostic data for iOS** IMPROVED 70

Updates Gen IR so the `--capture` flag now includes diagnostic data in `.xcarchive` files for iOS troubleshooting.

— Exact flag and file type named, directly usable Static Analysis updates (2026-09-03)

04 **veracode.json exclusion controls for Java** NEW 65

Adds `veracode.json` support for excluding specific Java `.jar` files from static analysis, and adds `veracode.json` support for excluding specific Java classes from

Static Analysis.

— Names exact config file and two exclusion targets Static Analysis updates (2026-09-03)

THINNER COVERAGE BELOW

05 **Expanded language, framework, and platform version support** NEW

55

Static Analysis adds support for numerous new language, framework, and platform versions: Dart 3.13/Flutter 3.47, Dart 3.12/Flutter 3.44, Dart 3.11/Flutter 3.41, Dart 3.10/Flutter 3.38, Dart 3.9/Flutter 3.35, Dart 3.8/Flutter 3.32, PHP 8.4 and PHP 8.5, Android 17, Next.js 16.x and 15.x, .NET 10, Kotlin 2.4 and 2.3, the `CloudKit` framework and Apple Security framework APIs for Apple Platforms, Xcode 26.4, JDK 26 and JDK 25 (LTS), scanning and viewing Java source code on the platform, React 19, Angular 20, Go 1.26 and Go 1.25, Apex 64 and Apex 65, Spring Boot 4/Spring Core 7/Spring Security 7, the Svelte framework, Ruby 4.0 and Rails 8.1, Swift standard library taint propagation on iOS, the Go Fiber web framework, Micronaut 4.x, Node.js 22.x, taint propagation for Apple Platforms framework APIs, initial iOS 26 support, Fast HTTP in Go, and Azure Functions v4 on .NET 9 in isolated process mode.

— Long name list but no mechanism beyond added coverage

Static Analysis updates (2026-09-03)

06 **Hardcoded secrets detection in Java config files** IMPROVED

50

Expands hardcoded secrets detection to include application configuration files with `-local` and `-prod` naming patterns in Java.

— Names the pattern but no scan-side action given Static Analysis updates (2026-09-03)

07 **Expanded third-party library detection** IMPROVED

50

Expands third-party library detection to include Fluxnova in Java, and expands support for `pnpm` modules in JavaScript to improve dependency resolution.

— Names the library and package manager but little mechanism

Static Analysis updates (2026-09-03)

08 **Pipeline Scan support for artifacts over 200 MB** IMPROVED

45

Adds support for scanning packaged artifacts larger than 200 MB in Pipeline Scan.

— States the limit but no configuration step given Static Analysis updates (2026-09-03)

09 **Automatic copybook detection in COBOL parsing** IMPROVED

25

Enhances COBOL parsing to automatically detect copybooks.

— Bare statement with no further detail Static Analysis updates (2026-09-03)

WAS THIS USEFUL? 👍 👎

OX Security

Feedback ⓘ 2 RELEASES · 2026-09-16 BLOG ↗ RANK

Accelerate the transition to software that secures itself. OX Security pinpoints risk to the exact line of code, eliminating security debt from AI generation to cloud runtime.

> OX Security launched OX Cloud, a unified CNAPP platform bundling CSPM, KSPM, DSPM, runtime vulnerability detection, cloud inventory and attack-path mapping, alongside new AIDR and AASM capabilities for detecting and governing AI agents, plus reachability-based prioritization to strip unreachable findings from the backlog.

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

01 **OX Cloud CNAPP platform launch** NEW 75

OX Cloud launches as a unified CNAPP covering `CSPM` (continuous detection of cloud misconfigurations and compliance enforcement across all cloud accounts), `KSPM` (Kubernetes security risk and policy-compliance monitoring), `DSPM` (discovery, classification and protection of sensitive data across cloud environments), Runtime Vulnerabilities (identifying vulnerabilities actively exposed in running workloads and prioritizing what is truly exploitable), Cloud Inventory (a continuously updated inventory of every workload, identity, data store, Kubernetes cluster and AI agent running in the environment, including shadow AI and unmanaged non-human identities), Cloud Graph and Attack Path Mapping (visualizing relationships across cloud assets including lateral movement paths and blast radius), and Runtime Events (investigating runtime anomalies with full context to accelerate incident detection and triage).

— Names each CNAPP module and its scope but no config or API surface.

OX Cloud: CNAPP Coverage Plus Runtime S… OX Security Launches OX Cloud, Redefini…

02 **AIDR: AI Detection and Response** NEW 70

AIDR (AI Detection and Response) detects and governs AI-driven usage, configuration and threats across cloud and runtime environments in real time, covering AI agents calling tools, models accessing data, and MCP servers executing autonomously. Alert detail views show who or what acted, what they touched and what else they could have reached.

— Describes real-time detection mechanism and alert view but no API/config.

OX Cloud: CNAPP Coverage Plus Runtime S… OX Security Launches OX Cloud, Redefini…

03 **AASM: Agentic Attack Surface Management** NEW 65

Agentic Attack Surface Management (AASM), powered by the OX Agent, discovers AI agents, models, MCP servers and non-human identities and maps what they can reach, uncovering hidden backdoors and reachable exposures through agentic-level inspection.

— Explains discovery mechanism but no configurable surface given.

OX Cloud: CNAPP Coverage Plus Runtime S… | OX Security Launches OX Cloud, Redefini…

THINNER COVERAGE BELOW

04 **Reachability-based risk prioritization** NEW 55

Reachability-based prioritization filters out risks that cannot actually be reached in cloud environments, eliminating unreachable misconfigurations, vulnerable packages and supply-chain exposures from the findings backlog so that only exploitable, evidenced risks surface.

— Explains filtering logic but gives no metrics or config detail.

OX Cloud: CNAPP Coverage Plus Runtime S… | OX Security Launches OX Cloud, Redefini…

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Dalfox

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗ CODE ↗ RANK

Dalfox is a powerful open-source XSS scanner and utility focused on automation.

> Dalfox v3.2.3 expands its scanning surface to GraphQL and XML/SOAP bodies, deepens DOM-XSS detection, adds a new filter-bypass technique, and reworks scan concurrency and cross-origin safety for faster, safer scans.

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

01 **Sub-not-gsub filter bypass detection** NEW 78

Dalfox detects the sub-not-gsub angle filter bypass: a doubled `<<svg ...>>` payload opens a real tag where a one-shot `str::replace` only consumed the first `<`.

— Names exact payload and underlying replace mechanism v3.2.3

02 **Faster concurrent parameter scanning** IMPROVED 70

Scans now run concurrently within a parameter, ramp payload batches instead of jumping straight to a full `--workers` window, and exit early for endpoints that uniformly escape their echo, significantly cutting scan time.

— Describes mechanism and names the `--workers` flag v3.2.3

03 **Cross-origin form probing disabled** IMPROVED 68

Cross-origin `<form action>` targets are no longer probed during parameter discovery, preventing `-H` headers and `--cookies` from being sent to page-controlled origins.

— Names exact flags affected and the security behavior fixed v3.2.3

04 **Wider DOM-XSS AST coverage** IMPROVED 65

DOM-XSS static analysis now covers Promise combinators, optional chaining, and static bracket paths, with more sinks and sources recognized and fewer false positives reported.

— Explains scope and effect but no runnable flag given v3.2.3

THINNER COVERAGE BELOW

05 **GraphQL and XML/SOAP injection surfaces** NEW 50

Dalfox now treats GraphQL and XML / SOAP request bodies as new injection surfaces for XSS scanning, alongside its existing HTTP parameter targets.

— Names the surfaces but no flag or endpoint to invoke them v3.2.3

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Google Cloud Forensics Utils

Feedback ⓘ

1 RELEASE · 2026-09-16 | CODE ↗ | RANK

Python library to carry out DFIR analysis on the Cloud

> Cloud Forensics Utils added zone control for image exports and extended container explorer support in imaging instances.

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

01 Zone parameter for ExportImage NEW

80

The `ExportImage` method now accepts a `zone` parameter (default `'us-central1-a'`), letting callers specify the zone used for Cloud Build imaging instances instead of always using the default region.

Export a GCP compute image to a specific zone rather than accepting the default us-central1-a — useful when your GCS bucket and build resources are in a different region.

python

```
image.ExportImage(
    gcs_output_folder='gs://my-forensics-bucket/exports/',
    image_format='qcow2',
    output_name='disk-export',
    zone='europe-west1-b'
)
```

— Named parameter, default value, and runnable example given 20260916

THINNER COVERAGE BELOW

02 Expanded container explorer dependencies IMPROVED

45

Adds `libbtrfs-dev` and `libgpgme-dev` to the forensics packages startup script, expanding container explorer support in imaging instances.

— Names packages but no usage detail or mechanism 20260916

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ 1 RELEASE · 2026-09-10

BLOG ↗ RANK

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

> Finite State expanded firmware unpacking to a new device model, tightened SBOM generation-context metadata, and improved detection of static-library components and Windows executable resources.

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

01 **Static library component detection and scanning** IMPROVED

70

Additional static-library Components are now identified where there was previously no signature coverage, including `picolibc`, `PCRE2`, `glib`, and LLVM runtimes. Static library archives (`.a` files) are now treated as a single Component and scanned directly, instead of being unpacked into hundreds of object-file rows.

— Names four libraries and the `.a` archive handling change.

September 10, 2026 - Firmware unpacking…

02 **SBOM generation context marked after-build** IMPROVED

60

Binary scan SBOMs now record generation context as `after-build`, matching the CISA SBOM minimum elements.

— Names the exact field value and compliance reference.

September 10, 2026 - Firmware unpacking…

THINNER COVERAGE BELOW

03 **Clearer SPDX 3.x import error** IMPROVED

55

Importing an SPDX 3.x SBOM now returns a clear error stating that SPDX 3.x is not yet supported, instead of a generic format-detection failure.

— Names the format and the improved error behavior.

September 10, 2026 - Firmware unpacking…

04 **Phoenix Contact FL SWITCH 2000 firmware unpacking** NEW

50

Finite State now unpacks Phoenix Contact FL SWITCH 2000 firmware updates so the Linux root filesystem and its Components appear in scan results, producing a complete

SBOM for the device.

— Names the device and SBOM outcome but no unpacking mechanism.

September 10, 2026 - Firmware unpacking…

05 **Windows executable embedded resources retained** IMPROVED

45

Embedded resources in Windows executables, including those used by Delphi applications, are now kept and analyzed instead of being discarded.

— States the before/after but no scan-level detail.

September 10, 2026 - Firmware unpacking…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sublime Security

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗

RANK

Sublime Security is a cloud email security platform that detects and remediates advanced threats using AI-powered analysis.

> Sublime Security added the ability to ingest Microsoft 365 Junk reports as user-reported phishing signal, without requiring end-users to change their reporting workflow.

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

01 **Microsoft 365 Junk report ingestion** NEW 67

A new opt-in setting under `Admin > Account > Abuse Mailbox` lets Sublime ingest Outlook `Junk` reports as user reports for Microsoft 365 message sources, capturing phishing signal without requiring end-users to use the Report Phishing button.

— Names exact setting location and behavior, but no limits or scope given.

Platform v1.88.228 (Microsoft365 Junk R…

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

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

> Fastly launched AI Runtime Control (ARC), a new control plane and firewall for governing LLM traffic at the edge, and extended its analytics APIs to track ARC traffic volume.

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

01 **arc_requests metric in analytics APIs** NEW 60

Fastly added the `arc_requests` metric to the Historical Stats and Real-Time Analytics APIs to track requests received by AI Runtime Control.

— Names exact metric and both APIs, but no usage example.

Add AI Runtime Control requests metric …

THINNER COVERAGE BELOW

02 **AI Firewall and AI Runtime Control (ARC)** NEW 53

Fastly added AI Runtime Control (ARC), a control plane that sits between applications and LLM providers to govern AI traffic, along with Fastly AI Firewall, a security module for ARC that inspects large language model (LLM) requests.

— Names the two components and their role but no config or setup steps. product docs

WAS THIS USEFUL? 👍 👎

◆ IDS / IPS

OISF Suricata

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗ CODE ↗

RANK

Suricata is a network Intrusion Detection System, Intrusion Prevention System and Network Security Monitoring engine developed by the OISF and the Suricata community.

> Suricata 8.0.7 expands the firewall policy engine with new ARP and default-action controls, adds subdomain matching to datasets, introduces a unified TCP session-matching keyword, lets app-layer-protocol match multiple protocols per rule, and brings firewall progress states to the SMTP parser.

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

01 Firewall policy engine ARP and default-action options NEW

75

Adds `accept-arp` option under `firewall.policies` in `suricata.yaml` to control whether ARP traffic is accepted by the firewall policy engine, and adds `default-policy` setting to specify a fallback action when no explicit policy matches any rule.

Allow ARP traffic through the firewall and set a default drop policy for unmatched flows.

yaml

```yaml
firewall:
  policies:
    accept-arp: true
    default-policy: drop
```

— Names config keys and file with a runnable example suricata-8.0.7

02 Subdomain matching in dataset keyword NEW

65

Adds `match subdomain` option to the `dataset` keyword, enabling rules to match subdomains of entries in a dataset, such as detecting connections to any subdomain of a blocklisted domain.

Detect connections to any subdomain of a blocklisted domain set using the new `match subdomain` dataset option.

```
$ alert dns any any -> any any (msg:"Blocklisted domain hit";
dns.query; dataset:isset,bad-domains,type string,match
subdomain; sid:1000002; rev:1;)
```

— Named keyword option with a working rule example `suricata-8.0.7`

03 **Multi-protocol matching in app-layer-protocol** IMPROVED

65

Extends the `app-layer-protocol` keyword to accept a pipe-separated list of protocol values in a single rule, allowing one rule to match multiple protocols (e.g. `app-layer-protocol:dns|http`).

Match DNS or HTTP traffic in a single rule using the new pipe-separated protocol list in `app-layer-protocol`.

```
$ alert tcp any any -> any any (msg:"DNS or HTTP detected"; app-layer-protocol:dns|http; sid:1000001; rev:1;)
```

— Named keyword with concrete example syntax `suricata-8.0.7`

THINNER COVERAGE BELOW

04 **tcp.session keyword for TCP lifecycle matching** NEW

40

Adds `tcp.session` keyword for unified TCP lifecycle matching across a session.

— Names the keyword but no example or mechanism detail `suricata-8.0.7`

05 **Firewall progress states in SMTP parser** NEW

40

Adds firewall progress states to the SMTP parser, enabling policy enforcement at SMTP transaction stages.

— States the addition but no example or stage detail `suricata-8.0.7`

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ 2 RELEASES · 2026-09-16 CHANGELOG ↗

RANK

Sysdig Secure is a container security platform providing runtime threat detection, forensics, and compliance monitoring for cloud-native environments.

> Sysdig Secure reduced KSPM scan overhead and significantly expanded the CLI Scanner with standalone binary detection, binary allow/deny policies, license reporting, and a new OS distribution.

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

01 Per-package license reporting in CLI Scanner NEW 95

CLI Scanner v1.30.0 adds per-package license data (including license category) to the JSON scan result under the v1 schema, available via `--output json`, in Technical Preview. It also adds `--show-licenses` to print the ten most frequent licenses and their occurrence counts in console output, and `--full-licenses-table` to show every license found, both in Technical Preview.

Audit an image for license compliance by getting a full per-package license breakdown alongside vulnerability results in machine-readable JSON.

```
$ sysdig-cli-scanner --output json <image>
```

Quickly review the most common open-source licenses present in an image during a CI pipeline scan.

```
$ sysdig-cli-scanner --show-licenses <image>
```

— Multiple named flags with runnable commands and schema detail.

Secure SaaS - September 16, 2026

02 Standalone binary scanning in CLI Scanner NEW 85

Sysdig CLI Scanner v1.30.0 adds standalone binary scanning via `--detect-unknown-executables` to identify and report vulnerabilities in applications not installed by a package manager, including executables the scanner does not recognize.

Scan a container image and surface vulnerabilities in binaries that were not installed via a package manager, including unrecognized executables.

```
$ sysdig-cli-scanner --detect-unknown-executables <image>
```

— Named flag plus runnable command demonstrating exact usage.

Secure SaaS - September 16, 2026

THINNER COVERAGE BELOW

03 **CIS Linux Benchmark filesystem scan efficiency in KSPM Analyzer**

IMPROVED

50

KSPM Analyzer v1.49.5 improves CIS Linux Benchmark filesystem scans by excluding kubelet pod volume mounts, reducing IOPS on nodes.

— Names mechanism and benefit but no config surface to act on.

changelog-20260916-4e5c3a19

04 **Binary Allow List and Deny List rules in CLI Scanner** NEW

45

CLI Scanner v1.30.0 adds support for Binary Allow List and Binary Deny List rules, so vulnerability policies that allow or deny specific binaries are now evaluated by the CLI Scanner.

— Names the policy types but no flag or command shown. Secure SaaS - September 16, 2026

05 **Hummingbird OS distribution support** NEW

25

CLI Scanner v1.30.0 adds support for the Hummingbird OS distribution.

— Bare mention of new OS support with no further detail. Secure SaaS - September 16, 2026

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ 1 RELEASE · 2026-09-16 NOTES ↗ RANK

Plerion is a cloud security platform that discovers and manages vulnerabilities and misconfigurations across infrastructure.

> Plerion's Terraform onboarding now supports all three workload scanning deployment models, including a fully self-managed service account setup that keeps scanning infrastructure inside the customer's own organization.

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

01 **Service account model for Terraform onboarding** NEW

82

Terraform onboarding now supports all three workload scanning deployment models, including a service account you own and Plerion operates, keeping scanning inside your organization without appliances in scanned accounts. Two Terraform templates provision this setup: one grants Plerion the required access, and one deploys per-region resources (appliance network, auto-scaling scan pool, and snapshot encryption key); both register with Plerion automatically during `terraform apply` with no manual console steps.

— Names deployment models, two templates, resources, and terraform apply command.

Onboard with Terraform and scan in a se…

WAS THIS USEFUL? 👍 👎

Upwind

Feedback ⓘ 1 RELEASE · 2026-09-14 BLOG ↗ RANK

Upwind is a runtime-powered CNAPP that protects cloud workloads by combining posture, detection and response, and identity security with real-time runtime context.

> Upwind Code now covers all four major version control providers by adding Azure DevOps and Bitbucket Cloud integrations, plus a generic API path for other providers and a new console page to manage scan scope.

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

01 **Azure DevOps and Bitbucket Cloud integrations in Upwind Code**

NEW 70

Upwind Code adds an Azure DevOps integration that connects at the organization level through Microsoft Entra ID to discover repositories and surface findings on pull requests without modifying existing CI/CD pipelines, and a Bitbucket Cloud integration that connects at the workspace level to automatically include new repositories and surface findings within pull request review workflows. Together with existing GitHub and GitLab support, this completes coverage across all four major version control providers.

Connect an Azure DevOps or Bitbucket Cloud environment to begin repository discovery and pull request scanning.

> 📍 In the Upwind console, go to Components → Integrations and select Azure DevOps or Bitbucket Cloud to connect your organization or workspace.

— Names auth method and connection level but no API/CLI detail

Upwind Code Expands Enterprise Coverage…

02 **Organizations & Repositories management page** NEW 60

Adds an Organizations & Repositories page in the Upwind console to centrally manage which organizations and repositories are scanned, keeping access and scan scope separate.

Control which organizations and repositories are actively scanned after connecting a version control provider.

> 📍 In the Upwind console, go to Organizations & Repositories to select which connected organizations and repositories are in scope for scanning.

— Gives exact UI navigation to control scan scope

Upwind Code Expands Enterprise Coverage…

THINNER COVERAGE BELOW

03 **API reporting for unsupported version control providers** NEW

30

Teams using a version control provider other than GitHub, GitLab, Azure DevOps, or Bitbucket Cloud can report scan results to Upwind through the API.

— Mentions an API but gives no endpoint or method

Upwind Code Expands Enterprise Coverage…

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Chainguard

Feedback ⓘ 1 RELEASE · 2026-09-14 NOTES ↗ RANK

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

> Chainguard added 26 new container images this week, many with FIPS variants, and announced a breaking default-user change for the sonar-scanner-cli image.

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

01 Sonar-scanner-cli image switches to nonroot default

BREAKING

95

Effective 2026-10-12, the `sonar-scanner-cli` image default user changes from root (UID 0) to nonroot (UID 65532), and the scanner scratch directory moves from `<project>/.scannerwork` to `/tmp/.scannerwork`, which breaks pipelines that read `report-task.txt` from the old location and derived images that write to system paths as root. To restore previous behavior, run with `--user 0` (Docker), `USER root` (Dockerfile), or `runAsUser: 0` (Kubernetes), or set `-Dsonar.working.directory=<project>/.scannerwork` and ensure the workspace is writable by the container user.

– Exact flags, paths, UIDs and migration commands provided Week of 2026-09-14

02 26 new container images added, including FIPS variants

NEW

60

Chainguard added new container images: `clamav-prometheus-exporter` (application + FIPS), `crossplane-azure-marketplaceordering` (application), `kserve-sklearnserver` (ai), `commercial-spilo-fips` (FIPS), `graphql-hive-gateway` (application), `jupyterhub-k8s-singleuser-sample` (application + FIPS), `nuxeo-2025` (application), `chainguard-server-hypervisor-aws` (base), `mongodb-kubernetes` (application + FIPS), `opentelemetry-python-instrumentation-fips` (FIPS), `strimzi-kafka-fips` (FIPS), `crossplane-aws-emrcontainers` (application + FIPS), `kro` (application + FIPS), `pyrra` (application + FIPS), `starboard-exporter` (application + FIPS), `krakend` (application + FIPS), and `mongodb-search` (application + FIPS).

– Names every image and type but gives no pull command or registry path

Week of 2026-09-14

└──▷ BREAKING ON UPGRADE

! Effective 2026-10-12: the `sonar-scanner-cli` image default user changes from root (UID 0) to nonroot (UID 65532); the scanner scratch directory moves from `<project>/.scannerwork` to `/tmp/.scannerwork`, breaking pipelines that read `report-task.txt` from the old location

and derived images that write to system paths as root. To restore previous behavior, run with `--user 0` (Docker), `USER root` (Dockerfile), or `runAsUser: 0` (Kubernetes), or set `-Dsonar.working.directory=<project>/.scannerwork` and ensure the workspace is writable by the container user.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Managed Detection & Response (MDR)

Arctic Wolf

Feedback ⓘ 1 RELEASE · SEEN 2026-09-16 PRESS ↗ RANK

Instinct meets AI. Arctic Wolf Aurora's built-in agents operate at machine speed while human experts maintain oversight.

> Arctic Wolf launched Aurora MDR Connect, a new MSP-exclusive MDR tier offering agentless deployment and broad telemetry ingestion across the attack surface.

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

01 **Aurora MDR Connect MSP tier launch** NEW 45

A new MSP-exclusive MDR tier with agentless deployment that ingests telemetry from endpoint, network, cloud, identity, SaaS applications, and other sources to deliver detection and response across a broader attack surface than endpoint-only approaches.

— Names data sources and deployment model but no config or API detail.

launch-20260916-6b0c85b8

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Sophos Protection for Linux

Feedback ⓘ

2 RELEASES · SEEN 2026-09-16 NOTES ↗ RANK

Sophos Protection for Linux is an endpoint security solution providing antivirus, anti-malware, and threat detection for Linux servers and workstations.

> Sophos Protection for Linux 2026.2 and 2026.3 add tamper resiliency and Update Cache/Message Relay roles for the SPL Agent, a new diagnostic CLI tool, expanded Sophos Data Lake telemetry, and support for newer Linux kernels.

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

01 **Endpoint Self Help diagnostic CLI tool** NEW 65

Adds an Endpoint Self Help command-line tool for administrators to view product health and collect diagnostic information from a Linux server.

– Names the exact CLI tool and its purpose, but no usage syntax given Version 2026.3New

02 **Expanded telemetry events in Sophos Data Lake** NEW 60

Expands the `account_activity_linux` data lake event to include account creation, deletion, and modifications to group membership. Adds DNS activity and Server Name Indication (SNI) hostname events to the Sophos Data Lake for expanded network visibility, and adds container information to supported detection events in the Sophos Data Lake.

– Names each event type added but no query or access instructions Version 2026.3New Version 2026.2

THINNER COVERAGE BELOW

03 **Update Cache and Message Relay roles for SPL Agent** NEW 55

Enables Linux servers running SPL Agent to be configured as Update Caches and Message Relays in Sophos Central.

– Describes the new role but not configuration steps Version 2026.3New

04 **Tamper Resiliency for SPL Agent files** NEW 40

Adds automatic detection and recovery for unauthorized changes to SPL Agent files (Tamper Resiliency).

– Explains the mechanism at a high level, no configuration detail Version 2026.3New

05 **Data lake processing performance and batching improvements**

IMPROVED

30

Improves event processing performance for endpoints with heavy container activity, and introduces data-lake batching updates to increase alert fidelity.

— Describes improvement direction only, no metrics or mechanism Version 2026.2

06 **Expanded Linux kernel support** NEW

30

Adds support for kernel 6.19/7.0, and later adds support for kernel 7.1.

— Names kernel versions but no other detail Version 2026.3New Version 2026.2

WAS THIS USEFUL? 👍 👎

Automox

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

Check out our full list of release notes, and stay up to date with what

> Automox improved Remote Tools' Splashtop Streamer auto-install, adding automatic reinstall on re-imaged devices and automatic retry of transient failed installs.

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

01 **Splashtop Streamer auto-install improvements** IMPROVED

55

Re-imaged devices with auto-install enabled now have the Splashtop Streamer reinstalled automatically without manual intervention after an imaging batch. Failed installs now retry automatically when the failure is transient (network interruption, locked file, or antivirus interference), while permanently unresolvable failures (e.g. unsupported OS) are reported immediately.

— Explains mechanism and failure handling but no config surface named

Remote Tools — Splashtop Streamer auto-…

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ 5 RELEASES · 2026-09-15 → 2026-09-16

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's headline change this window is centralizing pattern update scheduling under Version Control Policies ahead of the 2027 End of Life for Standard Endpoint Protection and Server & Workload Protection; it also added Outlook-based phishing simulation reporting, OpenLDAP identity sync for Security Awareness, and incremental hardening of Service Gateway firmware and the Private Access Connector.

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

01 Pattern update scheduling in Version Control Policies NEW

91

Version Control Policies gains pattern update schedule configuration, reachable via Endpoint Security → Endpoint Security Configuration → Version Control Policies, for agents on version 202609 and later. Engine updates now follow agent software update settings when managed through Version Control Policies, consolidating previously separate controls, and enrollment into Version Control Policies now automatically enrolls all unenrolled Protection Manager instances in a single operation. Agents older than version 202609 continue to follow pattern update settings from Protection Manager rather than Version Control Policies, splitting the management path for mixed fleets.

Centralize pattern update scheduling for all Protection Manager instances before Standard Endpoint Protection and Server & Workload Protection reach End of Life in 2027.

> 📍 In the console, go to Endpoint Security → Endpoint Security Configuration → Version Control Policies, open the enrollment workflow, and confirm enrollment to automatically enroll all unenrolled Protection Manager instances.

— Exact console path, agent version cutoff, and migration caveat given

Version Control Policies now supports m…

02 Outlook Add-in phishing simulation reporting NEW

77

Trend Vision One adds a Report Phishing button via the Outlook Add-in so users can report suspicious or simulated phishing emails directly from Microsoft Outlook. Reports feed simulation analytics through a new response option added to the phishing results filter, a new user status in simulation profiles records the method used to report a phishing email, and risk scores are reduced for users who demonstrate consistent, proactive phishing reporting behavior over time.

Enable the Outlook Add-in so end users can report phishing simulation emails directly from their inbox, feeding results into simulation analytics and risk scoring.

> 📍 In the console, go to Cyber Risk Exposure Management › Security Awareness › Phishing Simulations, then enable the Add-in for Outlook.

— Names button, filter option, and enablement path

Integrate Email Security Outlook Add-in…

---

03 **Service Gateway firmware 3.0.34 platform updates** NEW

70

Service Gateway firmware 3.0.34 adds CLI commands to enable and disable MicroK8s debug logging, adds package signature verification before firmware upgrade and service installation, adds Brazil (BR) site FQDN support, extends custom CA support to the `sg-apl` service and Node.js-based services, upgrades the appliance OS to Rocky Linux 9.8, and enhances diagnostics collection with MicroK8s inspection reports and container image checks.

— Lists every firmware change but no exact CLI syntax

Service Gateway firmware version 3.0.34…

THINNER COVERAGE BELOW

04 **OpenLDAP identity source in Security Awareness** NEW

57

Security Awareness adds OpenLDAP as a supported identity data source, enabling direct user import and synchronization from LDAP-based directories for phishing simulations and training campaigns.

Connect an OpenLDAP directory so users and groups are automatically available for targeting in phishing simulations and training campaigns, replacing manual CSV uploads.

> 📍 In the console, go to Cyber Risk Exposure Management → Security Awareness → Phishing Simulations and add OpenLDAP as an identity data source, then select users or groups to include in a campaign.

— Names integration and console path, no field-level detail

Open LDAP support added as an identity …

---

05 **Diagnostics and telemetry improvements in Private Access Connector**

IMPROVED

37

Zero Trust Secure Access Private Access Connector 3.9.0.2162 improves application access visibility with richer telemetry data in Secure Access history and adds more detailed diagnostic logging for Browser Access to aid troubleshooting.

— Describes improvements without naming fields or logs

Zero Trust Secure Access Private Access…

└──▷ BREAKING ON UPGRADE

! Agents older than version 202609 continue to follow pattern update settings from Protection Manager, not from Version Control Policies — mixed fleets will have split management paths after this change.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

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

> Google SecOps launches Google Unified Security, bundling Chrome Enterprise Premium integration, multi-cloud CDIR via Security Command Center Enterprise, and AI Protection/Model Armor into the platform, alongside a new Grok `match_all` parser option and two upcoming breaking changes affecting OAuth scopes and legacy IOC feeds.

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

01 Multi-cloud CDIR via Security Command Center Enterprise NEW

90

Google SecOps adds Cloud Detection, Investigation, and Response (CDIR) via Security Command Center Enterprise (SCC-E) integration, covering Google Cloud, AWS, and Azure within a single converged platform. Cloud Asset Inventory collection flows into the Google SecOps Entity Graph, enriching alerts with multi-cloud asset context (region, owner, metadata tags); ingestion of AWS CloudTrail and GuardDuty events and Azure AD activity logs enables multi-cloud SIEM log searching using natural language, Yara-L, or raw log syntax; and CDIR Curated Detections for Cloud Threat scenarios include enhanced privilege escalation alerts that automatically map to Google SecOps cases.

Activate multi-cloud CDIR in Google SecOps to surface Security Command Center threat alerts and ingest AWS and Azure logs for unified investigation.

> 📍 1. Activate the Enterprise Tier of SCC-E using your existing Google SecOps instance. 2. Validate Security Command Center Detection Services are enabled. 3. Connect Security Command Center to AWS for Entity Graph resource data. 4. Ingest Google Cloud logs via direct ingestion and set up the Google Cloud Export Filter. 5. Ingest AWS logs into Google SecOps. 6. Ingest Azure logs into Google SecOps. 7. Enable the CDIR Curated Detections.

— Names SCC-E, log sources, and gives full activation steps product docs

02 Chrome Enterprise Premium integration in SecOps NEW

85

Google SecOps adds Chrome Enterprise Premium (CEP) integration with curated detections enriched with Safe Browsing and Google Threat Intelligence (GTI) for browser-based threats such as malicious extensions, credential theft, and data exfiltration. It also adds dedicated dashboards for Chrome security events covering

managed and BYOD devices, and automated response playbooks for Chrome Enterprise actions including blocking malicious URLs, blocking extensions, and quarantining devices.

Onboard Chrome Enterprise Premium to Google SecOps to immediately gain browser threat visibility and enable curated detections enriched with GTI.

> 📍 1. Configure the Chrome Enterprise Premium Reporting Connector to send data to Google SecOps. 2. Connect Chrome Enterprise Premium to Google SecOps. 3. Enable Chrome Enterprise Threats Curated Detections.

— Names sources, detection types, and playbook actions with setup steps `product docs`

---

03 **chronicle.readonly OAuth scope loses write access** BREAKING

80

Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow performing write operations must be updated to use the `chronicle` OAuth scope instead.

— Exact scopes and date given but no migration command `changelog-20260916-a3061ae6`

---

04 **Legacy IOC feeds deprecated in favor of GTI** DEPRECATED

70

The `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed after March 18, 2027; affected feeds must be migrated to Google Threat Intelligence (GTI) before that date.

— Names exact feeds and deadline but no migration steps `changelog-20260916-a3061ae6`

---

05 **AI Protection and Model Armor for AI workloads** NEW

65

Google SecOps adds AI Protection within Security Command Center Enterprise, providing discovery and cataloging of AI assets (models, datasets, applications), virtual red-teaming to surface toxic attack paths, and detection and response for LLM hijacking, data poisoning, and AI-targeted threats, with response workflows integrated into Google SecOps. Model Armor protects AI workloads against data loss, brand damage, and offensive content across models and cloud environments.

— Describes mechanism and scope but gives no setup steps `product docs`

---

06 **match_all option in Grok parser filter** NEW

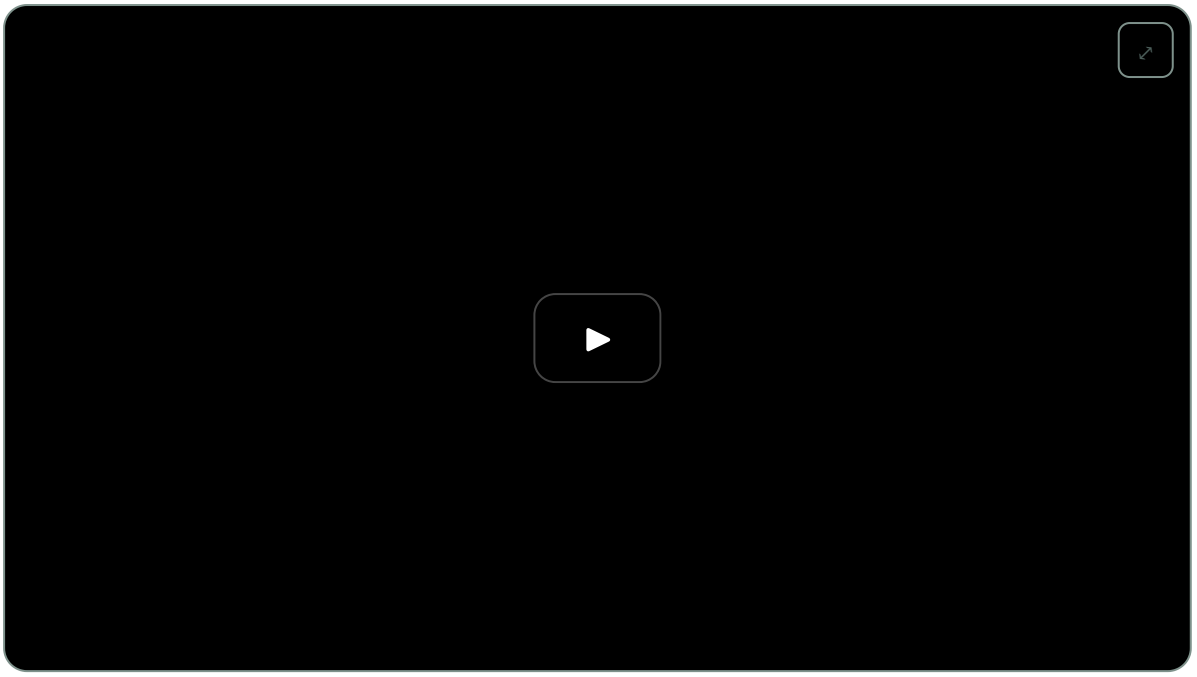
65

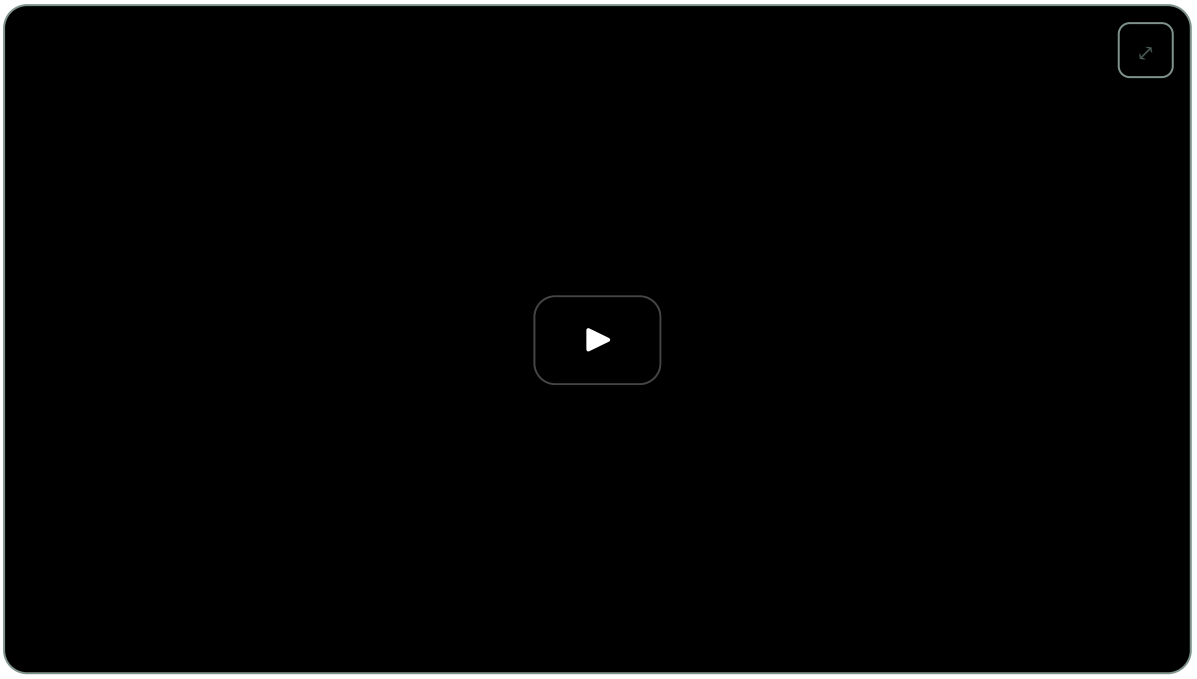
Google SecOps parser syntax adds a `match_all` option to the Grok filter, letting parsers extract all non-overlapping pattern occurrences within a field instead of only the first match.

— Names the exact filter option but no usage example `September 15, 2026 (2026-09-15)`

changelog-20260916-a3061ae6

└──▷ ALSO FROM THESE RELEASES





└──▷ BREAKING ON UPGRADE

! Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow performing write operations must be updated to use the `chronicle` OAuth scope instead.

! The `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed after March 18, 2027; affected feeds must be migrated to GTI before that date.

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ 5 RELEASES · 2026-09-10 → 2026-09-15 BLOG ↗

RANK

Databricks is a unified analytics platform for data engineering, data science, and machine learning built on Apache Spark.

> Databricks' biggest move this window is taking OpenSharing GA for zero-copy sharing of foreign Delta tables, alongside a beta for managed tables in Lakeflow pipelines and betas extending ABAC and Data Classification to Unity Catalog views.

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

01 **Managed tables and append flows in Lakeflow pipelines** NEW

90

Lakeflow pipelines now support creating managed tables and writing to them via one or more append flows using `append_flow` (Beta). `import_checkpoint` on `append_flow` migrates existing Structured Streaming workloads into a Lakeflow pipeline without reprocessing the source, and `create_table` plus `CREATE TABLE ... FLOW` syntax define managed tables and flows within pipelines.

— Names exact syntax and migration mechanism Managed tables in Lakeflow pipelines (B…

02 **READ METADATA privilege for compliance security profile workspaces**

NEW

75

The `READ METADATA` privilege will be available by default for workspaces with the compliance security profile enabled, granting read-only visibility into security-sensitive metadata — including privilege grants, row filters, column masks, and ABAC policies — that is not exposed by `BROWSE`.

.
.
.

— Explains exact scope versus BROWSE, no exact steps

The READ METADATA privilege will soon b…

03 **Sharing foreign Delta tables via OpenSharing** NEW

60

OpenSharing now generally available for sharing foreign Delta tables from OneLake, Hive metastore, or AWS Glue without copying data into Databricks.

— Names sources and GA status but no command shown

Sharing foreign Delta tables with OpenS…

04 **ABAC and Data Classification extended to Unity Catalog views**

NEW

60

ABAC row filters and column masks now extend to views in addition to tables (Beta, DBR 19+), letting existing Unity Catalog policies automatically apply to views once the `ABAC on Views` preview is enabled from the account console's Previews page. Data Classification separately gains the ability to scan Unity Catalog views (Beta), classifying columns for sensitive data using the same detection engine used for tables.

Enable the ABAC on Views beta so that existing ABAC policies automatically apply to views, not just tables.

> 📍 In the account console, go to Previews and enable `ABAC on Views`.

— Gives preview toggle location but limited mechanism detail

`ABAC on views is in Beta (2026-09-15)` `Data Classification now scans Unity Cat…`

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Qualys

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

Cybersecurity blog covering software, services, and risk insights. Stay updated with the latest on threat detection, compliance, and vulnerability management.

> Qualys TruRisk Eliminate gains AI-powered patch reliability scoring and automated, EDR-free host isolation for overdue CISA KEV vulnerabilities, plus phased ring-based patch rollout.

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

01 Automated host isolation via Cloud Agent NEW

90

TruRisk Eliminate can now automatically isolate hosts natively through the Qualys Cloud Agent when a user-defined QQL policy is met, with no EDR integration required; the host is quarantined within seconds and a timestamped audit record captures which rule, host, and deadline triggered the action. A built-in CISA KEV isolation use case applies the QQL `vulnerabilities.status:[ACTIVE,NEW,REOPENED] AND vulnerabilities.riskFactor.cisaKEVDueDate < "now-1s" AND vulnerabilities.riskFactor.cisaKnownExploits:TRUE` to isolate hosts where a KEV remediation deadline has passed and the vulnerability remains open.

Automatically isolate any host where a CISA KEV remediation deadline has passed and the vulnerability is still open, preventing exploitation while patching continues.

```
$ vulnerabilities.status:[ACTIVE,NEW,REOPENED] AND
vulnerabilities.riskFactor.cisaKEVDueDate < "now-1s" AND
vulnerabilities.riskFactor.cisaKnownExploits:TRUE
```

— Provides exact QQL policy syntax and isolation mechanism. 2026-09-15-2b585d21

02 AI-powered Patch Reliability Scoring in Patch Management NEW

86

TruRisk Eliminate adds AI-Powered Patch Reliability Scoring that combines LLM-analyzed public sentiment with Qualys telemetry on patch rollback and vulnerability reopen rates to produce a pre-deployment reliability score (Low / High) for each patch. A new `patch.kb` QQL filter in the Patch Management app (Patches > Windows) surfaces these assessments for specific KBs; clicking `Patch Reliability: Low` on a result opens an AI Assessment window showing known issues and source links via `Detailed Report`.

Check patch reliability before deploying a Patch Tuesday update — filter to specific KBs in the Patch Management app to see whether Qualys rated them Low or High reliability before pushing to production.

> 📍 In the Patch Management application, go to Patches › Windows and enter the QQL query: patch.kb: [KB5124008 , KB5124012] — then click `Patch Reliability: Low` on any result to open the AI Assessment window and review known issues and source links via `Detailed Report` .

— Names exact QQL filter, UI path, and scoring inputs. `2026-09-15-2b585d21`

THINNER COVERAGE BELOW

03 **Ring-based phased patch deployment** NEW 35

TruRisk Eliminate Patch Management supports ring-based phased deployment, allowing low-reliability patches to be progressively validated through Test → Staging → Production environments before broader rollout.

— Describes rollout stages but no config or command detail. `2026-09-15-2b585d21`

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices on networks without credentials.

> runZero added API support for suppressing and unsuppressing vulnerabilities, vulnerability groups, and findings.

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

01 API endpoints to suppress vulnerabilities and findings NEW

45

New API endpoints let users suppress and unsuppress vulnerabilities, vulnerability groups, and findings.

— Names the capability but no endpoint paths or methods given

runZero 5.1.260915.0 (2026-09-15)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Netskope

Feedback ⓘ 1 RELEASE · SEEN 2026-09-16 PRESS ↗ RANK

Netskope is a cloud security platform that monitors and protects data across cloud applications, web services, and private apps.

> Netskope introduced Agent Action Control, a new capability in its AI Security suite that lets security teams inspect and block risky AI agent actions before they execute.

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

01 **Agent Action Control for AI agents** NEW 40

A new addition to the Netskope AI Security suite, Agent Action Control enables security teams to inspect and block risky AI agent actions before they execute.

— Names the feature but no config, UI path, or mechanism detail. launch-20260916-9bf6fac9

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Cisco Duo Desktop

Feedback ⓘ 1 RELEASE · SEEN 2026-09-16

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's latest updates add macOS 27 support alongside new telemetry, environment, and device-posture reporting capabilities in version 7.21.

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

01 **Optional telemetry consent for usage data** NEW 60

Duo Desktop now prompts users at startup to opt in to error and usage data collection. This can be changed at any time in app preferences, or administrators can set it system-wide to bypass the per-user prompt entirely.

Suppress the telemetry consent prompt for all users on managed endpoints by setting the system-wide preference before rollout.

> 📍 In the app's preferences, navigate to the telemetry/data collection setting and enable or disable it system-wide to bypass the per-user startup prompt.

– Describes mechanism and admin override, but no config key named `snapshot-20260916`

THINNER COVERAGE BELOW

02 **Virtual environment detection for GCE and Azure** NEW 30

Duo Desktop can now detect Google Compute Engine and Microsoft Azure as virtual environments.

– Names two providers but no mechanism or user action `snapshot-20260916`

03 **OS update status and management agent reporting** NEW 25

Duo Desktop now reports OS automatic update status and management agent detection for Secure Access.

– States the capability but no detail on how it's surfaced `snapshot-20260916`

04 **macOS 27 (Golden Gate) support** NEW 25

Duo Desktop adds support for macOS 27 (Golden Gate).

– Bare compatibility statement with no further detail `product docs`

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback

3 RELEASES · 2026-09-15 → 2026-09-16

CHANGELOG ↗ RANK

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

> DNSFilter shipped a GA Model Context Protocol server for connecting AI assistants to DNS and threat data, a GA Unblock Request Management workflow with CSV export, a redesigned MSP Branding page, and iOS Roaming Client v1.4.5 with several MDM configuration and device-sync fixes.

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

01 **Unblock Request Management (GA)** NEW 80

Adds the Unblock Requests page to the dashboard, providing a full request queue with Allow or Deny actions, threat-category warnings on risky domains, an Investigate Mode row action that scopes directly to that domain in the Query Log, and CSV export of active and historical requests.

Review and act on pending user unblock requests, flagging risky domains before approving — useful during a morning triage workflow.

> 📍 In the dashboard, go to Unblock Requests › select a pending request › review any threat-category warnings › click Allow or Deny. To investigate traffic context, use the Investigate Mode row action to scope directly to that domain in the Query Log. Export the full queue via the CSV export button for audit records.

— Names the page, every action and the export control, no API endpoint.

09/15 September 15 – Web Application an…

02 **MDM key behavior fixes in iOS Roaming Client v1.4.5** IMPROVED 80

Three MDM configuration keys got behavior fixes in iOS Roaming Client v1.4.5: `dns_over_tls_enabled` now accepts `false` from first device registration and persists the setting across DNS extension restarts; changes to `dhcp_dns_search_domains` now take effect within one configuration refresh with no DNS extension restart required; and the `userName` MDM field now populates the Last Logged in User field immediately on initial device registration.

Disable DNS over TLS from your MDM payload so the setting is enforced from the moment a device first registers.

```yaml
dns_over_tls_enabled: false
userName: jane.doe@example.com
dhcp_dns_search_domains:
  - example.com
  - corp.example.com
```

— Names three config keys with before/after behavior and a config example.

09/15 September 15 - iOS Roaming Client…

03 **DNSFilter MCP Server for AI assistants** NEW 75

GA launch of the DNSFilter MCP Server under Integrations → AI Connectors, connecting AI assistants (Claude, ChatGPT, Codex) to DNS activity, threat data, policy management, and allow/block lists via the Model Context Protocol. Users copy the server URL from the AI Connectors card and paste it into their AI client as the MCP server endpoint to complete authentication.

Connect an AI assistant to DNSFilter for natural-language policy queries and threat investigation via the MCP Server.

> 📍 In the DNSFilter dashboard, go to Integrations → AI Connectors, open the DNSFilter MCP Server card, copy the server URL, then paste it into your AI client (Claude, ChatGPT, or Codex) as the MCP server endpoint and complete authentication.

— Names integration path and protocol but no API/endpoint schema given.

changelog-20260916-e4ed8a25

04 **Branding page replaces Whitelabel settings** IMPROVED 65

Redesigns the Whitelabel settings page as a dedicated Branding page, exposing Logo, Favicon, Dashboard Name, and Subdomain URL as individual, first-class branding objects. New MSPs provisioned after September 15 get a branded subdomain (`yourdomain.dnsfilter.com`) configured with current branding defaults.

— Names each branding field but no steps for migrating existing MSPs.

09/15 September 15 – Web Application an…

05 **Device registration and sync improvements in iOS Roaming Client**

IMPROVED 65

The `.mobileconfig` file now bundles DNSFilter root certificates alongside the iOS Roaming Client configuration so a single MDM profile push installs both; reinstalling the client now reuses the existing device record instead of creating a duplicate registration

entry; and registered devices now sync with the DNSFilter dashboard only on device state changes instead of every 60 seconds.

— Names the config file and sync interval but changes are mostly automatic.

changelog-20260916-e4ed8a25 | 09/15 September 15 - iOS Roaming Client…

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone's macOS/iOS client added support for device trust backed by X.509 certificates.

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

01 **X.509 certificate-backed device trust** NEW 25

The macOS/iOS client now supports device trust backed by X.509 certificates.

— No detail on setup, enrollment, or configuration mechanism.

macOS / iOS 1.5.20 (2026-09-15)

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

JumpCloud

Feedback ⓘ 5 RELEASES · 2026-09-16 BLOG ↗ RANK

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

> JumpCloud's biggest addition this window is Agent Identities (Public Preview), a non-human identity system for AI agents with OAuth credentials, group-based SSO access, and MCP server authorization; alongside it JumpCloud shipped in-place Private Repository app updates, expanded SCIM group mapping, several MFA and access-risk-detection features, breaking changes to the reactivate-user endpoint and Dynamic Groups API filters, and a round of Password Manager desktop app improvements.

└──▷ WHAT SHIPPED · 23 FEATURES most completely described first ? what's the number?

01 Agent Identities for non-human AI agent management (Public Preview)

NEW 95

Adds Agent Identities (Public Preview) under `Identity Management > Agents` for managing non-human identities such as AI agents: OAuth client credentials are generated on the agent's Authentication tab for the JumpCloud AI Gateway, group-based SSO application access is granted via `Identity Management > Agent Groups`, and MCP server authorization is configured from the agent's Applications tab via Manage connections.

Register and configure an AI agent identity so its access to SSO applications is governed independently from human user accounts.

📍 1. Go to Identity Management › Agents and click + Add. 2. Enter the agent name, owner, and company email on the Details tab. 3. On the Authentication tab, click Generate to create OAuth client credentials for the JumpCloud AI Gateway. 4. Go to Identity Management › Agent Groups, create a group, add the agent, and assign the target SSO applications. 5. On the agent's Applications tab, click Manage connections to authorize MCP servers on behalf of the agent.

— Names every menu path and credential mechanism with a runnable setup sequence.

Release Notes 2026 (2026-09-16)

02 Reactivate-user endpoint replaced

BREAKING 90

Deprecates `GET api/systemusers/reactivate/{id}`, which now returns a 404, in favor of `POST https://console.jumpcloud.com/systemusers/reactivate/{id}`

accepting `{ activation_email: {email_address} }` in the request body; any script or integration using the old endpoint must be updated.

Update an existing automation that reactivates a user account to use the new POST endpoint instead of the deprecated GET request.

```
$ curl -s -X POST
'https://console.jumpcloud.com/api/systemusers/reactivate/{id}'
-H 'x-api-key: <api_key>' -H 'Content-Type: application/json' -d
'{"activation_email": "user@example.com"}'
```

— Exact old and new endpoints, method, and request body given.

Release Notes 2024 (2026-09-16)

---

03 **Dynamic Groups API filter operator rename** BREAKING 80

Dynamic Groups API `memberQuery` filters now use `$lte` and `$gte` operators, replacing the deprecated `$le` and `$ge` syntax, aligning filter syntax across all API list endpoints; requests using the old operators must be updated.

— Exact old and new operator names given, migration action clear.

Release Notes 2024 (2026-09-16)

---

04 **Access Risk Detection with Slack alerting** NEW 80

Enable Access Risk Detection to monitor risk factors such as impossible travel, dormant accounts, and unfamiliar devices, tune detection thresholds, and route high-severity alerts to a configured Slack webhook based on minimum severity level.

Enable Access Risk Detection and route high-severity alerts to Slack so the security team is notified of impossible-travel or dormant-account anomalies in real time.

> 📍 1. Go to Configure Access Risk Detection in the Admin Portal. 2. Enable Access Risk Detection and select the risk factors to monitor (e.g., impossible travel, dormant accounts, unfamiliar devices). 3. Tune thresholds to match your environment. 4. Under Notifications, add a Slack webhook and set the minimum severity level that triggers it.

— Names risk factors and gives a full setup sequence with Slack routing.

Release Notes 2026 (2026-09-16)

---

05 **In-place app version updates in Private Repository** NEW 75

Adds an Update Version action on an application's Details tab in the Private Repository, letting admins upload a new binary for an existing custom app without

creating a new application, reassigning devices, or rebinding device groups; Windows EXE updates require entering an install success exit code before deployment.

— Names exact UI action and exit-code requirement but no runnable command.

Release Notes 2026 (2026-09-16)

06 **Modern TLS cipher enforcement for LDAP** BREAKING 75

Enforces modern TLS cipher suites for all LDAP connections, limiting to `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256`, removing CBC-mode and SHA-1/SHA-256 suites.

— Names every retained cipher suite exactly; no migration steps given.

Release Notes 2025 (2026-09-16)

07 **macOS Ventura update delay/block policy** NEW 70

New software management policy lets admins delay or block macOS Ventura upgrades fleet-wide until app compatibility is validated, configured via Device Management > Policies > New Policy > macOS > Software Update > Delay or Block macOS Ventura with a configurable deferral period bound to a target device group.

Block macOS Ventura upgrades fleet-wide until your org has validated app compatibility, using the new delay/block software management policy.

📍 In the Admin Portal, go to Device Management › Policies › New Policy › macOS › Software Update, select Delay or Block macOS Ventura, configure the deferral period, and bind the policy to your target device group.

— Exact navigation path and configuration steps given. Release Notes 2022 (2026-09-16)

08 **macOS Managed Login Items policy** NEW 70

New policy under Device Management > Policies > macOS > Managed Login Items lets admins add the JumpCloud Agent (or other persistent management software) to a managed items list, preventing macOS 13 Ventura end users from disabling it via System Settings.

Prevent end users on macOS 13 Ventura from toggling off the JumpCloud Agent or other persistent management software via System Settings.

📍 In the Admin Portal, go to Device Management › Policies › New Policy › macOS › Managed Login Items, add the JumpCloud Agent to the managed items list, and bind the policy to your macOS 13 device group.

— Exact navigation path and behaviour prevented are given.

Release Notes 2022 (2026-09-16)

09 **Escaped pipe characters in API list filters** IMPROVED 70

API list-endpoint filters now support a backslash escape for literal pipe characters inside `in` filter values, e.g. `filter=department:$in:Finance \| Accounting|IT`, previously impossible.

Query users whose department is `Finance | Accounting` or `IT` using a pipe literal inside an `in` filter — previously impossible without the new backslash escape.

```
$ curl -s -H 'x-api-key: <api_key>'
'https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT'
```

— Runnable curl example shows exact escape syntax. Release Notes 2024 (2026-09-16)

10 **Expanded SCIM group management** IMPROVED 70

Adds a dedicated mapping UI for prebuilt and custom SCIM apps supporting name overrides, extended schema mapping, and side-by-side object transformation previews; custom activation validation now supports selecting custom group names and representative test users and groups instead of hard-coded defaults, with all sync rules and activities logged to Directory Insights.

— Describes mechanism and logging but no UI path given. Release Notes 2026 (2026-09-16)

11 **Push Bombing Mitigation for MFA** NEW 65

Adds Push Bombing Mitigation settings under Security > MFA Configurations (JumpCloud Protect), including a sixty-second lockout window and configurable maximum concurrent push attempts, to counter fraudulent MFA push floods.

Enable Push Bombing Mitigation settings after a wave of fraudulent MFA push attempts — navigate to MFA Configurations to review or adjust the sixty-second lockout window.

> 📍 In the Admin Portal, go to Security › MFA Configurations, locate the JumpCloud Protect Push Bombing settings, and adjust the maximum concurrent attempts limit or disable the mitigation as needed.

— Names the setting and lockout window but not exact defaults for all fields.

Release Notes 2023 (2026-09-16)

12 **Global MFA requirement for MSP admins** NEW 65

Multi-Tenant Portal adds a Global MFA Requirement toggle under Account > Settings to enforce MFA for all MSP admins in a tenant.

Enforce MFA for all MSP admins in a tenant from the Multi-Tenant Portal to ensure no admin account is left unprotected.

📍 In the Multi-Tenant Portal, go to Account › Settings and enable the Global MFA Requirement toggle.

— Exact toggle name and navigation path given. Release Notes 2024 (2026-09-16)

13 **MDM Certificate & Token Status Board** NEW 60

New Multi-Tenant Portal Homepage widget shows MDM, SSO, and AzureAD certificate/token expiry status across all managed tenant organizations from a single dashboard.

Monitor MDM and SSO certificate expiry across all managed tenants from a single MSP dashboard to prevent auth disruptions.

📍 In the Multi-Tenant Portal, navigate to the Homepage and open the MDM Certificate & Token Status Board to review expiry status for MDM, SSO, and AzureAD tokens across all tenant organizations.

— Names the dashboard and what it monitors, with clear navigation.

Release Notes 2023 (2026-09-16)

14 **Webhooks for Access Requests** NEW 60

Access Requests now support configuring a webhook (e.g. pointing to a JIRA instance) under Access Requests > Settings so each new request submission automatically creates a ticket and notifies the team in real time.

Set up a webhook so Access Requests automatically create JIRA tickets and notify your team in real time when a user submits a request.

📍 In the Admin Portal, go to Access Requests › Settings, configure an approval flow assigning the relevant admins, then add a webhook pointing to your JIRA instance URL so each new request submission creates a ticket automatically.

— Clear setup path but no endpoint or payload format given.

Release Notes 2025 (2026-09-16)

15 **User group exclusion for Unified MFA** NEW 60

Adds a User Group Exclusion option under MFA > JumpCloud Protect > Unified MFA settings, letting admins exclude a group from Push MFA enforcement so those users

authenticate with TOTP only while remaining in the unified JumpCloud Protect experience.

Exclude a remote team group from Push MFA enforcement so those users authenticate with TOTP only, without losing the unified JumpCloud Protect experience.

> 📍 In the Admin Portal, go to MFA › JumpCloud Protect › Unified MFA settings, select User Group Exclusion, and add the target group to the exclusion list.

— Exact setting name and navigation path given. Release Notes 2025 (2026-09-16)

THINNER COVERAGE BELOW

16 **sso_auth log event for forced re-authentication** NEW

45

Adds `sso_auth` event to logs for forced re-authentication events, supporting compliance and security auditing.

— Names the exact event field but no query or usage detail.

Release Notes 2025 (2026-09-16)

17 **JumpCloud Go feature-flag removed for SaaS discovery** IMPROVED

40

Removes the requirement for JumpCloud Go to be enabled as a separate feature for discovering unmanaged applications in SaaS Management; the JumpCloud Go browser extension alone is now sufficient.

— States the change but no steps to verify or configure. Release Notes 2024 (2026-09-16)

18 **Security groups in Microsoft 365 directory integration** NEW

35

Adds M365/Azure AD security groups support to JumpCloud's Microsoft 365 directory integration.

— States what was added but no configuration detail. Release Notes 2022 (2026-09-16)

19 **Multiple IP addresses in RADIUS server config** IMPROVED

35

Adds multiple IP address support to RADIUS server configurations.

— One-line addition with no configuration detail. Release Notes 2024 (2026-09-16)

20 **Automatic termination of unresponsive Remote Assist agents** IMPROVED

35

Remote Assist now automatically detects and terminates unresponsive agent instances on end-user devices when an admin initiates a connection.

— Describes behaviour but no trigger conditions or thresholds.

Release Notes 2024 (2026-09-16)

21 **Asset cloning in Asset Management** NEW 35

Adds asset cloning capability in Asset Management to streamline creation of assets with shared hardware configurations.

— States capability without steps or mechanism. Release Notes 2025 (2026-09-16)

22 **Frequency-based alert triggers in Insights** NEW 35

Adds alert trigger conditions based on event frequency to Insights rules, so alerts fire only when defined criteria are met rather than on every single event.

— States the capability without configuration detail or thresholds.

Release Notes 2026 (2026-09-16)

23 **Settings tabs UI refresh** IMPROVED 25

Refreshes the Organization Profile, Security, and Features tabs in Settings with shorter labels and easier-to-scan helper text.

— Purely cosmetic description with no mechanism or scope. product docs

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; any script, automation, or integration using this request must be updated to use `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `{ activation_email: {email_address} }` in the body.

! Dynamic Groups API `memberQuery` filter operators `$le` and `$ge` are deprecated and replaced by `$lte` and `$gte`; API requests using the old operators must be updated.

! MCP servers added before this release without an associated SSO application or bookmark are now flagged in AI Gateway requiring an application assignment — access is not scoped by User Groups until one is assigned.

! The custom app creation option has been removed from the Add MCP Server dialog; a new application must now be created via the `+ New App` button which opens the standard Add New SSO Application wizard.

! The `.appimage` Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

WAS THIS USEFUL? 👍 👎

Aembit

Feedback ⓘ 1 RELEASE · 2026-09-15 BLOG ↗ RANK

Aembit is IAM for agentic AI and workloads. Enable secretless access and real-time policy enforcement across any environment to protect data and accelerate AI adoption.

> Aembit's sole release this window adds CrowdStrike AIDR as a new Content Security capability on Access Policies, giving admins inline allow/block/transform enforcement over MCP tool traffic tied to identity and policy context.

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

01 CrowdStrike AIDR Content Security on Access Policies NEW

88

Aembit adds Content Security as a new capability type on Aembit Access Policy objects, letting administrators attach a CrowdStrike AIDR content-security configuration to the same enforcement path that governs MCP access. It covers three MCP inspection checkpoints per policy — tool listings, tool inputs, and tool outputs — with AIDR returning an allow, block, or transform verdict that Aembit enforces inline. Every CrowdStrike AIDR finding is correlated with Aembit's identity and policy context (agent, target MCP server, applicable policy, Content Security decision) in a single workload event, and the integration is configurable through the Aembit UI, API, or Terraform with no changes required to existing AIDR console configuration.

Attach CrowdStrike AIDR content inspection to a specific MCP Access Policy so that tool listings, inputs, and outputs are evaluated before the agent's request is allowed through.

📍 In the Aembit UI, open the target Access Policy, navigate to the Content Security section, select CrowdStrike AIDR as the provider, and enable the checkpoints (tool listings, tool inputs, tool outputs) you want evaluated. Save the policy — AIDR verdicts (allow/block/transform) will now be enforced inline for every MCP connection governed by that policy.

— Names capability type, three checkpoints, verdicts, and UI/API/Terraform config paths.

CrowdStrike AIDR Tells You What's Risky…

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ 2 RELEASES · 2026-09-15 → 2026-09-16

CHANGELOG ↗ RANK

Read docs, explore use cases, learn best practices

> Ping Identity's Advanced Identity Cloud gained OAuth 2.0 Client ID Metadata Document (CIMD) support, Rich Authorization Details in access tokens, and new OAuth provider settings for audience token introspection and DCR software-statement auditing, while PingOne added custom client IDs for imported apps, AI agent detection in Bot Detection, and a modernized Audit UI, and DaVinci began enforcing a baseline content security policy on production flows.

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

01 **Token introspection by audience members** NEW 95

Adds the Allow audience members to introspect tokens setting on the OAuth 2.0 provider, permitting any client whose `client_id` appears in a token's `aud` claim to introspect that token without needing a dedicated scope. Enable it in the admin console via Native Consoles > Access Management > Realms > <realm> > Services > OAuth2 Provider > Advanced (Services > OAuth 2.0 Provider in Advanced Identity Cloud), toggling Allow audience members to introspect tokens .

Enable resource servers to introspect tokens they receive via token exchange, without needing a dedicated scope — configure the new provider setting in the OAuth 2.0 provider.

> 📍 In the admin console, go to Native Consoles › Access Management › Realms › <realm> › Services › OAuth2 Provider › Advanced, and enable Allow audience members to introspect tokens .

Enable resource servers to introspect tokens they receive via token exchange without requiring a dedicated scope — useful when deploying microservices that consume delegated tokens.

> 📍 In the Advanced Identity Cloud admin console, go to Native Consoles › Access Management › Services › OAuth 2.0 Provider, then enable the Allow audience members to introspect tokens setting and save.

— Exact setting name, mechanism, and console navigation given 14 Sep 2026 (2026-09-15) changelog-20260916-0043515d

02 **DCR software statement persistence** NEW 95

Adds the Store Software Statement setting on the OAuth 2.0 provider to persist the software statement provided during dynamic client registration (DCR) in the client

configuration and return it in DCR read responses, useful for auditing which software statement a dynamically registered client used. Enable it via Native Consoles > Access Management > Services > OAuth 2.0 Provider > Dynamic Registration > `Store Software Statement in OAuth2 Client`.

Preserve DCR software statements in client configs so they are returned in read responses — useful for auditing which software statement a dynamically registered client used.

> 📍 In the admin console, go to Native Consoles › Access Management › Realms › <realm> › Services › OAuth2 Provider › Dynamic Registration, and enable `Store Software Statement in OAuth2 Client`.

Preserve software statements across DCR lifecycle calls so downstream systems can audit the original client assertion on every read response.

> 📍 In the Advanced Identity Cloud admin console, go to Native Consoles › Access Management › Services › OAuth 2.0 Provider, then enable the `Store Software Statement in OAuth2 Client` setting and save.

— Exact setting name, purpose, and navigation path given `14 Sep 2026 (2026-09-15)` `changelog-20260916-0043515d`

---

03 **Configurable SAML authentication request storage** NEW

70

Adds an `Authentication Request Storage` option on the hosted IdP to configure where SAML authentication requests are stored, enabling concurrent SAML SSO attempts across multiple browser tabs or browsers to complete independently.

— Explains behavior change but omits exact UI navigation `14 Sep 2026 (2026-09-15)`

---

04 **Custom redirect URI for backchannel transactions** NEW

65

Adds `redirectUrlType` and `customRedirectUrl` parameters to the backchannel transaction initialize endpoint, enabling customization of the redirect URI returned.

— Names the parameters but not the full endpoint method or path `14 Sep 2026 (2026-09-15)`

---

05 **AI agent detection in Bot Detection** NEW

65

Enhances the PingOne Protect Bot Detection predictor with new rules to identify specific computer-using agent (CUA) types for more granular AI-driven activity monitoring, available in Protect (Signals) SDK 5.6.10 and later; workforce flows may also require PingID Device Trust agent version 1.2 or later. Ping Identity also released Protect (Signals) SDK version 5.6.11 for Web and PingID Device Trust agent version 1.2 with improved data collection.

— Names SDK/agent versions but gives no setup steps `product docs`

06 **OAuth 2.0 Client ID Metadata Documents (CIMD)** NEW 60

Supports OAuth 2.0 Client ID Metadata Documents (CIMD), allowing AI agents and external clients to identify themselves via a hosted metadata document without prior or dynamic client registration.

— Explains mechanism but gives no setup steps or endpoint 14 Sep 2026 (2026-09-15)

07 **Rich Authorization Details in access tokens** NEW 60

Enables issuance of access tokens carrying `authorization_details` through the authorization code and refresh token grants, when details were granted during authorization.

— Names the claim and grant types but no config example 14 Sep 2026 (2026-09-15)

08 **Custom client ID/secret for imported apps** NEW 60

Adds support via the PingOne API for setting a custom client ID and client secret when importing OIDC-based applications and custom resources from PingFederate or third-party solutions into PingOne.

— Names the API and source systems but not the field name product docs

THINNER COVERAGE BELOW

09 **Error logging for reCAPTCHA Enterprise node** IMPROVED 50

Adds error logging to the reCAPTCHA Enterprise node when its Google project, site key, or API key secret is misconfigured.

— Names the node and misconfig cases but no fix steps changelog-20260916-0043515d

10 **LDAP and Identity Store Decision node v2.0** IMPROVED 45

Adds v2.0 of the LDAP Decision node and the Identity Store Decision node with improved username handling.

— Names the nodes and version but not what changed in handling 14 Sep 2026 (2026-09-15)

11 **Backchannel transaction subject persistence** IMPROVED 45

Backchannel transactions now record the authenticated subject on completion even when no session is created, preserving any subject declared before authentication.

— Describes behavior change but no config or endpoint named 14 Sep 2026 (2026-09-15)

12 **Android device binding EC key attestation** NEW 45

Supports EC key attestation for Android device binding, including compatibility with newer Google hardware attestation root certificates.

— Names the mechanism but gives no configuration path changelog-20260916-0043515d

13 **Baseline content security policy for DaVinci flows** BREAKING 45

Enforces a baseline content security policy (CSP) on all flows running on a PingOne domain in production environments, rolled out in phases.

— States rollout scope but not the CSP directives or migration steps product docs

14 **Modernized Audit UI** IMPROVED 30

Updates the Audit UI with a modernized interface and a simplified time zone selector.

— Describes the change but not what it replaces or where product docs

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1 RELEASE · 2026-09-11 BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password SaaS Manager rolled out a public beta AI Assistant and gave admins finer control over ChatGPT usage tracking, alongside a new Google Workspace field mapping and clearer AI spend reporting.

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

01 **ChatGPT organization selection for AI usage tracking** NEW 55

Adds the ability to choose which ChatGPT organization is used for AI usage tracking when automatic detection selects the wrong organization, preventing authorization errors for customers with separate OpenAI Platform and ChatGPT organizations.

— Explains the problem it solves but not where to configure it

1Password SaaS Manager #1701 (2026-09-1…

02 **AI Assistant public beta in SaaS Manager** NEW 45

Adds AI Assistant (public beta), enabled by default for new SaaS Manager customers, with a per-customer disable control available.

— Names the feature and default/opt-out behavior but no UI path

1Password SaaS Manager #1701 (2026-09-1…

03 **Google Workspace select-type field mapping** NEW 40

Adds support for mapping select-type custom fields from Google Workspace to SaaS Manager.

— Names the integration surface but gives no setup steps

1Password SaaS Manager #1701 (2026-09-1…

04 **AI Spend and Consumption reporting improvements** IMPROVED 40

Improves AI Spend and Consumption charts and search to display friendly product display names instead of raw product IDs, and improves error messages to clarify when usage cannot be attributed to an individual user.

— Describes UI/UX changes with no actionable steps

1Password SaaS Manager #1701 (2026-09-1…

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1 RELEASE · 2026-09-14 BLOG ↗ RANK

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

> Runlayer's 1.28.428 release centers on infrastructure and safety controls — Terraform-managed agents, HMAC-signed webhooks, and ToolGuard scan budgets with fail-open/closed handling — alongside new connector integrations, unified event triggers, expanded browser auto-deploy, and richer usage and spend reporting.

└──▷ WHAT SHIPPED · 13 FEATURES most completely described first ? what's the number?

01 **Scan budgets in ToolGuard** NEW 90

Adds scan budgets to ToolGuard: configure latency and payload budgets with `fail-open` / `fail-closed` behavior, `allow` / `truncate` / `block` payload actions, and budget-exceeded events written to audit logs.

— Names every mode and action verbatim with clear config surface.

1.28.428 - Webhook signing, scan budget…

02 **Terraform provider for agent management** NEW 65

Adds `runlayer_agent` resource and data source to the Terraform provider so agents can be fully managed as infrastructure-as-code.

— Names exact Terraform resource but gives no config example.

1.28.428 - Webhook signing, scan budget…

03 **HMAC signing for agent webhooks** NEW 65

Adds HMAC signing for agent webhooks — create, rotate, and disable signing secrets from the webhook dialog, with a scheme picker and curl examples.

— Names UI flow and curl examples but not the curl command itself.

1.28.428 - Webhook signing, scan budget…

04 **AI Watch browser auto-deploy expansion** NEW 65

Enables Chrome, Edge, and Firefox auto-deploy on managed Windows devices via the AI Watch browser extension, with native identity hosts per browser, and extends auto-deploy to Microsoft Edge on managed macOS devices.

— Names browsers and platforms but no deployment steps.

1.28.428 - Webhook signing, scan budget…

05 **Work ledger for units-of-work pricing** NEW 65

Adds a work ledger for units-of-work pricing showing how each unit's cost was attributed, with branch checkouts steering attribution and Runlayer artifacts and skills counted as outputs with in-app preview.

— Explains attribution mechanism but no navigation or command.

1.28.428 - Webhook signing, scan budget…

06 **Google Workspace and Drive connector additions** NEW 60

Adds `get_table`, `delete_between_text`, and `delete_table` tools to the Google Workspace connector, and server-side PDF text extraction to the Google Drive connector.

— Names all new tool identifiers but no usage detail.

1.28.428 - Webhook signing, scan budget…

07 **Unified Connector Events for triggers** IMPROVED 60

Unifies event triggers into Connector Events with audience fan-out, an event volume signal, and chat-list visibility for event-triggered runs.

— Describes mechanism and scope but no config path.

1.28.428 - Webhook signing, scan budget…

THINNER COVERAGE BELOW

08 **Spend breakdowns by person, team, department** NEW 55

Adds spend breakdowns per person, team, and department — with clients drilling down to models — and browseable usage by current teams and departments.

— Describes drill-down mechanism but no UI path given.

1.28.428 - Webhook signing, scan budget…

09 **Dynamic plugin routing per client** NEW 45

Adds dynamic plugin routing per client, with Codex configured for plugin routing and org installs learning client-routed tools.

— Names Codex example but no config or command shown.

1.28.428 - Webhook signing, scan budget…

10 **Upstream progress notification relay** NEW 40

Relays upstream progress notifications through the proxy so long-running tools report progress to clients.

— Explains behavior but no protocol or config detail.

1.28.428 - Webhook signing, scan budget…

11 **Usage APIs by IdP groups and departments** NEW 30

Adds usage APIs by IdP groups and departments.

— Bare description with no endpoint or field names.

1.28.428 - Webhook signing, scan budget…

12 **X API MCP connector added** NEW 26

Adds X API MCP to the connector catalog.

— Bare name, no scope or capability described. 1.28.428 - Webhook signing, scan budget…

13 **Per-agent artifact limit removed** IMPROVED 25

Removes the per-agent artifact limit.

— States change with no prior limit value or scope.

1.28.428 - Webhook signing, scan budget…

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗

RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> AI SAFE² CLI 0.5.0 introduces `safe2 evidence diagnose` for evidence-bounded failure localization across agentic AI systems, backed by new `safe2.failure-source.v1` and `safe2.failure-diagnosis.v1` JSON contracts, plus scanner exclusions and a portable agent delivery workflow.

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

01 **safe2 evidence diagnose command** NEW 95

Adds `safe2 evidence diagnose` to rank evidence-supported failure locations across an identified agent system, with `--system-identity`, `--output`, `--card`, and `--strict` flags. `--strict` mode returns a nonzero exit code when evidence cannot support one strong or moderate leading candidate, and `--card` emits a Failure Localization Card as a human-readable Markdown decision document alongside the JSON diagnosis. Every diagnosis is bound to a complete AI SAFE² system identity manifest (harness, tool, policy, memory, evaluator identity) and classifies evidence into five levels — Strong, Moderate, Limited, Insufficient, and Contradicted — and is compatible with Codex, Claude Code, Hermes, OpenClaw, and CI systems without designating any single provider as source of truth.

— Exact command with all flags and mechanism described

2026-09-15_CLI_0.5.0_Find_Where_Agent_W…

02 **Failure diagnosis JSON contracts** NEW 75

Introduces `safe2.failure-source.v1`, a provider-neutral schema for supplying observations and competing failure candidates to the diagnostic engine, and `safe2.failure-diagnosis.v1`, which returns deterministic rankings, evidence assessments, ownership, and recommended actions from a diagnosis run.

— Named schema versions but no usage example given

2026-09-15_CLI_0.5.0_Find_Where_Agent_W…

THINNER COVERAGE BELOW

03 **Scanner exclusions for AI SAFE² artifacts** IMPROVED 40

Adds scanner workspace exclusions to avoid traversing generated AI SAFE² evidence directories, named virtual environments, and uv caches during repository scans.

— No specific config key or flag named `2026-09-15_CLI_0.5.0_Find_Where_Agent_W…`

04 **Portable agent delivery workflow** NEW 30

Introduces a portable agent delivery workflow that guides agents through isolation, building, proof, self-assessment, independent review, and shipping with receipts.

— Only described in prose, no commands or config `2026-09-15_CLI_0.5.0_Find_Where_Agent_W…`

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗

RANK

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

> Docker Sandboxes shipped v0.43.0 alongside four documentation updates this window, headlined by a tri-state shared-skills-store access control system (with a breaking rename of the `shareSkills` config key), scoped credential management including GitHub Packages npm injection and MCP gateway secret namespacing, overflow model routing, and expanded environment-file variables, plus a long tail of CLI, diagnostics, and lifecycle refinements.

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

01 **Credential scoping and injection controls** NEW 95

`sbx` credential commands gain a `--sandbox` flag to scope a custom secret to a specific sandbox instead of the global default, and `sbx secret rm --sandbox` now immediately revokes the removed credential from the sandbox proxy. A `--registry-auth-endpoint` flag lets sandboxes trust a separate authentication host for self-hosted registries (e.g. a self-hosted GitLab instance) using a split Registry v2 Bearer auth flow. A documented workflow using `sbx credentials add custom --host npm.pkg.github.com --env NODE_AUTH_TOKEN --command 'gh auth token'` plus a `.npmrc` entry referencing `${NODE_AUTH_TOKEN}` injects GitHub Packages credentials into sandboxes for installing private npm packages without exposing the real token.

Trust a split-auth self-hosted GitLab registry so the sandbox proxy can exchange your stored PAT for a registry token at the advertised realm.

```
$ sbx credentials add registry \
  --sandbox my-app \
  --registry registry.example.com \
  --username "$GITLAB_USER" \
  --password "$GITLAB_PAT" \
  --registry-auth-endpoint https://gitlab.example.com/jwt/auth
```

Inject a dynamically refreshed GitHub token so sandboxed agents can install private npm packages from GitHub Packages without embedding the real token.

```
$ sbx credentials add custom \
  --sandbox my-sandbox \
  --host npm.pkg.github.com \
  --env NODE_AUTH_TOKEN \
  --command 'gh auth token'
```

Wire up the sandbox project so npm resolves private packages through the credential proxy using the injected placeholder.

```ini
@my-org:registry=https://npm.pkg.github.com
//npm.pkg.github.com/:_authToken=${NODE_AUTH_TOKEN}
```

— Runnable commands and config for every named credential surface `v0.43.0`

---

02 **Tri-state shared skills store access control** BREAKING

90

Environment files gain a `skills` field (`off` , `readonly` , or `readwrite` , defaulting to the daemon default of `readonly` unless overridden by the organization) and a matching `--skills` flag on `sbx run` and `sbx create` to set the shared skill store mount mode per sandbox. A new `skills.defaultMode` setting, configurable via `sbx settings set` , controls the default mode for all future sandboxes, and `--no-share-skills` remains as a deprecated alias for `--skills=off` . This replaces the old `shareSkills` key in `sbxenv.yaml` , which is no longer read — the new `skills` key must be used instead.

Launch a sandbox completely isolated from the shared skill store to keep it outside the shared trust boundary.

```
$ sbx run --skills=off claude
```

Set the default skill store access mode to read-only for all future sandboxes so agents can consume skills but never mutate the shared store.

```
$ sbx settings set skills.defaultMode readonly
```

Grant the sandbox write access to the shared agent skills store when the agent needs to persist new skills across sessions.

```yaml
skills: readwrite
```

— Names every flag, key and the exact breaking rename with migration path `v0.43.0`

---

03 **Environment file path variables and sandbox name overrides** NEW

85

Environment files support `${{ env.fileDir }}`, expanding to the absolute path of the directory containing the file (useful for locating scripts or assets stored alongside it), and `${{ env.projectDir }}`, resolving to the absolute path of the project directory passed to the command; a shared `~/.sbxenv.yaml` can use `workspace: ${{ env.projectDir }}` to mount each project's own directory. All `sbx env` subcommands (`sbx env create`, `sbx env run`, `sbx env rm`) now accept a `--name` flag to override the sandbox name set in the file or auto-generated.

Override the sandbox name at runtime so the same environment file can spin up isolated sandboxes for different branches without editing the file.

```
$ sbx env create --name feature-branch-test
sbx env run --name feature-branch-test -- npm test
sbx env rm --name feature-branch-test
```

Use `${{ env.fileDir }}` in a shared environment file to run a setup script stored next to the file, regardless of which project directory invokes it.

```yaml
lifecycle:
  post_create:
    - ${{ env.fileDir }}/scripts/setup.sh
```

Make a shared user-level environment automatically mount each project's own directory when launching a sandbox from any project folder.

```yaml
workspace: ${{ env.projectDir }}
```

— All variables, flag and subcommands named with working examples `v0.43.0`

---

04 **MCP gateway secret namespacing** BREAKING

85

Credential names starting with `mcp:` are reserved for the MCP gateway — these secrets stay on the host and are never injected into sandboxes, letting the gateway authenticate MCP server connections on behalf of sandboxed agents. In v0.43.0 the naming convention for MCP OAuth client secrets changed from `mcp:<server>.client_secret` to `mcp:<server>:client_secret`; any secret stored under the old name must be re-set with `sbx secret set mcp:<server>:client_secret`.

— Explains mechanism and gives exact migration command for the rename `v0.43.0`

---

05 **MCP command headers and SSRF check control** NEW 75

`sbx mcp add` gains a `--header` flag to send custom request headers to remote MCP servers, with values substituted from the local secret vault so tokens aren't exposed in shell history. MCP authorization also gains a `--skip-ssrf-check` flag for private OAuth discovery.

Connect an MCP server that requires a custom Authorization header sourced from the local secret vault, keeping the token out of shell history.

```
$ sbx mcp add https://mcp.example.com --header 'Authorization:
Bearer ${{ secret.mcp_token }}'
```

— Named flags with a working example, limited to two additions `v0.43.0`

---

06 **Overflow and hosted model provider routing** NEW 75

`sbx run --model` gains `--overflow-provider` and `--overflow-model` flags to pair a local model with a hosted fallback for requests that exceed the local model's context limit. `sbx run --provider` now also accepts hosted models.dev providers, served through llmman.

Run a local model but automatically fall back to a hosted provider when a request exceeds the local model's context limit.

```
$ sbx run --model ollama/qwen2.5-coder --overflow-provider
openai --overflow-model gpt-4o
```

— Names both flags and the provider mechanism with a runnable example `v0.43.0`

---

07 **Time-based filtering for sandbox queries, pruning and listing**

NEW 73

Sandbox queries gain an `until` filter accepting a duration, RFC 3339 timestamp, or Unix timestamp to scope results by when a sandbox stopped, alongside the existing `since=<duration>` filter. `sbx prune` gains the same Docker-style `until` filtering, and `sbx ls --json` output now includes a last-used timestamp field.

Restart a previously stopped local sandbox by name to resume work with its persisted files intact.

```
$ sbx run --name my-sandbox
```

— Names the filter formats and every command it now applies to v0.43.0

THINNER COVERAGE BELOW

08 **sbx diagnose checks mkfs.erofs compatibility** NEW 55

`sbx diagnose` now checks whether `mkfs.erofs` is usable and warns if its default block size exceeds the sandbox kernel's page size.

— Names the exact check and condition it warns on v0.43.0

09 **Windows connectivity to sandboxd via AF_UNIX sockets** NEW 45

Windows clients can now connect to `sandboxd` through filesystem `AF_UNIX` sockets.

— Names the daemon and socket type but no setup steps v0.43.0

10 **Faster sbx version invocation** IMPROVED 45

Plain `sbx version` invocations now return embedded version information without full CLI startup.

— Names the exact command and behavior change v0.43.0

11 **Minimum sandbox memory reduced to 512 MiB** IMPROVED 40

The minimum sandbox memory has been decreased to 512 MiB, suitable for shell use cases.

— Gives the exact new limit but no configuration detail v0.43.0

12 **Kit installs through registry mirrors with explicit port** IMPROVED 35

Kits can now be installed through registry mirrors configured with an explicit port.

— One-line addition with no example or config key shown v0.43.0

13 **Mount information in inspect output** IMPROVED 30

`daemon inspect` and `inspect` now show mount information for sandboxes.

— Single-line description with no further mechanism given v0.43.0

14 **Shallow Git repository support in clone-mode sandboxes** IMPROVED 30

Clone-mode sandboxes now support shallow Git repositories.

— Bare one-line addition with no usage detail v0.43.0

└──▷ BREAKING ON UPGRADE

! `shareSkills` in `sbxenv.yaml` is replaced by `skills`; the new key accepts `off|readonly|readwrite` and the old key is no longer read.

! MCP OAuth client secrets are renamed from `mcp:<server>.client_secret` to `mcp:<server>:client_secret`; any secret stored under the old name must be re-set with `sbx secret set mcp:<server>:client_secret`.

WAS THIS USEFUL? 👍 👎

WitnessAI

Feedback ⓘ 1 RELEASE · 2026-09-13 BLOG ↗ RANK

WitnessAI sits in the network path and monitors, controls and protects employee, application and agent AI activity.

> WitnessAI shipped a full new AI governance suite in this window: network-level discovery of AI tools and MCP connections, agentic tool-call governance, runtime prompt/response guardrails, four-tier policy enforcement with model routing, intent-based ML classification, and data tokenization, backed by auditable records for compliance reviews.

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

01 **Agentic Control for tool-call governance** NEW 75

Adds Agentic Control to enforce approved-tool and MCP allow-lists, block unauthorized tool calls before execution, perform pre-execution checks on approved tool calls, apply runtime controls to agent activity, and maintain a full audit trail of governed agent and tool-call activity across more than 100 LLM types.

— Describes mechanism and scale (100+ LLM types) but no config key

AI data classification: how it works (2…

02 **Four-tier policy control with model routing** NEW 75

Adds Control capabilities to enforce four-tier intelligent policies — allow, warn, block, and route — on AI traffic routed through the platform, enabling proportionate enforcement over prompt, upload, response, and tool-call interactions. The route action redirects sensitive queries (such as non-public drug research or financial data) from external models to approved internal models based on classified intent.

— Names all four policy tiers and a concrete routing example

AI data classification: how it works (2…

03 **Observe module for AI tool discovery** NEW 65

Adds Observe module for network-level discovery of AI tools and MCP connections — including native apps like Windows Copilot and Microsoft 365 — cataloging thousands of AI applications without endpoint clients or browser extensions, and attributing captured agent actions to a human identity.

— Names concrete apps and mechanism but no config or command surface

AI data classification: how it works (2…

04 **Data tokenization and detokenization** NEW 60

Adds data tokenization that replaces classified sensitive values (such as SSNs) before prompts leave the enterprise and policy-enforced detokenization that restores original

values for authorized users in model responses.

– Clear before/after mechanism but no config key shown

AI data classification: how it works (2…

THINNER COVERAGE BELOW

05 **Protect runtime guardrails for prompts and outputs** NEW

55

Adds Protect runtime guardrails to inspect prompts and filter harmful or off-policy responses in AI interactions routed through the platform, covering both inputs and generated outputs for PII and sensitive data per NIST Generative AI Profile guidance.

– Names NIST framework but no configurable rule surface

AI data classification: how it works (2…

06 **Intent-based machine learning classification** NEW

50

Adds intent-based machine learning classification that evaluates conversational context and purpose to detect sensitive data, suspicious behavior, and policy violations that keyword and regex matching cannot represent, with confidence scores adjusted by contextual signals.

– Explains mechanism but no operational surface to act on

AI data classification: how it works (2…

07 **Auditable records for AI risk reviews** NEW

50

Produces auditable records of what employees and agents sent to which model and what came back, providing evidence for AI risk reviews under frameworks such as DORA and the EU AI Act.

– Names compliance frameworks but no export format or endpoint

AI data classification: how it works (2…

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

GitHub CLI

Feedback ⓘ 1 RELEASE · 2026-09-15 NOTES ↗ RANK

GitHub's official command line tool

> GitHub CLI's device-code login flow now copies the OAuth code to the clipboard automatically, cutting a manual step from authentication.

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

01 Automatic clipboard copy for OAuth device codes IMPROVED

80

`gh auth login` and `gh auth refresh` now copy OAuth device codes to the clipboard by default, removing a manual copy step during authentication. This can be turned off with `gh config set clipboard disabled`.

Persistently opt out of automatic clipboard copying of OAuth device codes across all future auth flows.

```
$ gh config set clipboard disabled
```

— Names exact commands and opt-out config key, minor scope v2.101.0

WAS THIS USEFUL? 👍 👎

Syslifters SysReptor

Feedback ⓘ 1 RELEASE · 2026-09-16

NOTES ↗ RANK

A customizable and powerful penetration testing reporting platform for offensive security professionals. Simplify, customize, and automate your pentest reports with ease.

> SysReptor 2026.75 introduces a structured `json` field type for reports, turns on the AI agent by default, and fixes Jira project pagination alongside several smaller UI and workflow improvements.

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

01 **Automatic post-update tasks via update.sh** IMPROVED 60

Runs post-update tasks automatically after updating SysReptor via `update.sh`, replacing any manual post-upgrade steps.

— Names the exact script and removes manual steps 2026.75

THINNER COVERAGE BELOW

02 **Jira plugin pagination for project listing** IMPROVED 50

The `jira` plugin now follows API pagination to list all available projects, removing the previous cap on visible Jira projects.

— Explains mechanism and named plugin but no config detail 2026.75

03 **JSON field type for report fields** NEW 45

Adds a `json` field type for report fields, enabling structured data entry in findings and templates.

— Names the field type but not how to configure it 2026.75

04 **Image uploads in shared Excalidraw notes** NEW 35

Supports image uploads in shared Excalidraw notes, allowing diagrams to include raster images in collaborative sessions.

— Describes feature but not upload mechanism or limits 2026.75

05 **Improved project search result ranking** IMPROVED 35

Ranks project search results to prioritize name and tag matches over content matches.

— Explains ranking logic but no UI or config path 2026.75

06 **AI agent enabled by default** IMPROVED 30

Enables the AI agent by default in settings, so new deployments get AI-assisted reporting without manual activation.

— No named setting or flag given 2026.75

07 **Color indicators for retest statuses** IMPROVED 30

Adds color indicators for all retest statuses in the UI, making pass/fail/pending retest outcomes visually distinct at a glance.

— UI change described only in general terms 2026.75

08 **OffSec OSAI certification report design** NEW 20

Adds a built-in design for the OffSec OSAI certification report.

— Bare mention of a new report template 2026.75

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-09-16 CHANGELOG ↗ RANK

Connect, protect, and build everywhere.

> Cloudflare shipped a wide-ranging changelog: Workflows gained real-time event streaming, AI Gateway added BYOK-only enforcement, Workers got per-resource role-based access controls, and smaller additions touched DNS shadow-record detection, voice metrics, Python Hyperdrive, Access for Infrastructure tagging, Browser Run guardrails, and a DLP Passive Detection dashboard.

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

01 **Event streaming for Workflow instances** NEW 90

Adds WorkflowInstance.subscribe() and the `GET /subscribe` API endpoint to stream Workflow instance events (steps, attempts, sleeps, waits, rollbacks) without polling. Supports a `filter` parameter to narrow to specific event types and a `cursor` parameter to resume from a specific event.

— Full endpoint, method and parameters named, no example given

changelog-20260916-a29c4695

02 **BYOK enforcement for AI Gateway requests** NEW 90

Adds `byok_only: true` on the `PUT` gateway API request to require provider credentials for AI Gateway third-party requests and block fallback to Unified Billing, plus a per-request `cf-aig-no-wholesale: true` header for request-level enforcement.

Enforce BYOK credentials for all AI Gateway third-party provider requests and block fallback to Cloudflare-managed billing.

```
$ curl -X PUT
https://api.cloudflare.com/client/v4/accounts/<account_id>/ai-
gateway/gateways/<gateway_id> \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"byok_only": true}'
```

— Field, header and runnable curl example all provided changelog-20260916-a29c4695

03 **Per-turn metrics in @cloudflare/voice** NEW 80

Adds a `VoiceTurnMetrics` typed summary to `@cloudflare/voice` v0.4.0 via addEventListener('turnmetrics', ...), exposing per-turn `turnId`, terminal outcomes (`completed`, `no_output`, `output_limit`, `content_filtered`, `model_error`,

`tts_error`, `aborted`), and stage timings (speech-to-final-transcript, model-to-first-text, TTS-to-first-audio, total turn duration).

— Named API, version and fields; no usage example changelog-20260916-a29c4695

04 **Shadow DNS record detection in API and dashboard** NEW

70

Adds `include_shadow_metadata=true` query parameter to the DNS records API to return shadow metadata identifying delegating `NS` records and glue `A` / `AAAA` records for shadowed DNS entries. The Cloudflare dashboard now also surfaces shadowed record warnings on all zones, flagging DNS records that will not be served due to subdomain delegation.

— Query parameter named; dashboard surface described but no example

changelog-20260916-a29c4695

05 **Role-based access controls per Worker** NEW

70

Adds per-Worker role-based access controls with four roles — Metadata Read-Only, Content Read-Only, Editor, and Admin — configurable via the Cloudflare dashboard, API, or Terraform for teammates, agents, and CI/CD workflows.

— Four roles named; no exact API field or config path given changelog-20260916-a29c4695

06 **Tag-based target criteria in Access for Infrastructure** NEW

70

Adds tagged targets and tag-based target criteria (`include`, `require`, `exclude` operators) to Access for Infrastructure, manageable inline or via the Resource Tagging API.

— Operators and API named; no endpoint path given changelog-20260916-a29c4695

THINNER COVERAGE BELOW

07 **Hyperdrive support for Python Workers** NEW

40

Extends Hyperdrive support to Python Workers, enabling connections to PostgreSQL and MySQL databases.

— Thin description, no config or API details given changelog-20260916-a29c4695

08 **Passive Detection dashboard for DLP** NEW

25

Adds a Passive Detection dashboard showing detection totals, data type distribution, policy coverage, and detection entries.

— Only a dashboard screenshot description, no mechanism given changelog-20260916-a29c4695

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 16, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue