THE CYBER TOOLCHAIN

NO. 014

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 13, 2026 · EVERY DAY

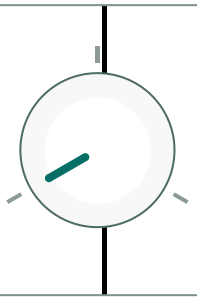
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 101 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ snapshot-20260713 NOTES ↗

Horizon3 NodeZero is an autonomous penetration testing platform that continuously finds and exploits real attack paths in your environment to identify and prioritize critical vulnerabilities.

NodeZero June 2026: Rapid Response overhaul, Real-Time View redesign, and External Pentest rebuild with reusable Scopes

└──▷ HOW TO FIND IT

When an alert fires during an active pentest, use the redesigned Real-Time View timeline and event-log filters to quickly confirm whether activity is NodeZero-generated or a genuine threat actor.

> 📍 In the NodeZero Portal, open an active pentest and navigate to Real-Time View, then use the timeline scrubber and event-detail filters to isolate events by source and timestamp.

Adds asset-level exploitability status, personalized remediation guidance, and verification feedback to Rapid Response, reducing vulnerability noise.

Redesigns Real-Time View (RTV) with a visual timeline, event-log summaries, and granular event details to help distinguish NodeZero testing activity from real attacks.

Rebuilds External Pentesting to replace Asset Groups with reusable Scopes, with improved bulk asset management, uploads, and support for large external attack surfaces.

Expands Rapid Response coverage to include Starlette/LiteLLM BadHost, Apache Solr, Palo Alto Networks PAN-OS, Nginx UI, Ivanti Sentry, Check Point Remote Access VPN, Splunk Enterprise, Cisco Unified Communications Manager, and Squid.

Adds MSP license-management improvements enabling parent accounts to manage clients' product access, asset allocations, and custom branding more easily.

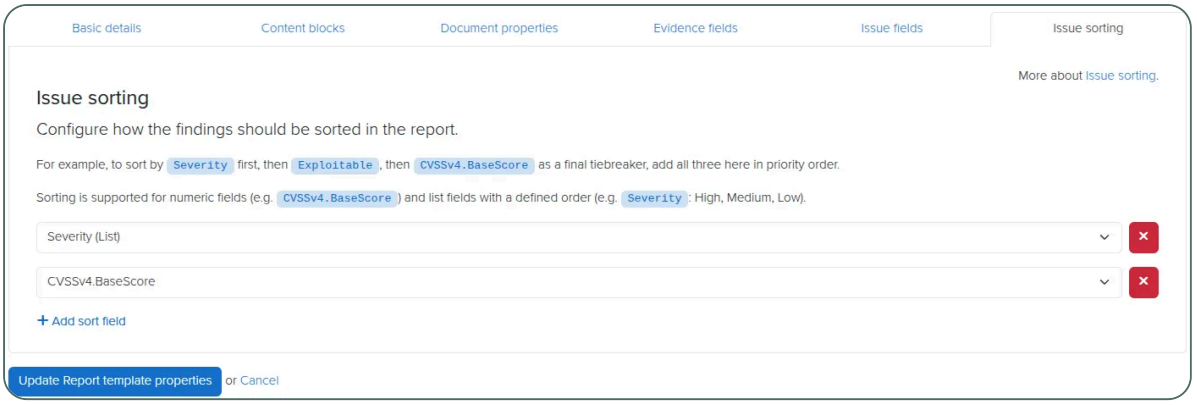
WAS THIS USEFUL? 👍 👎

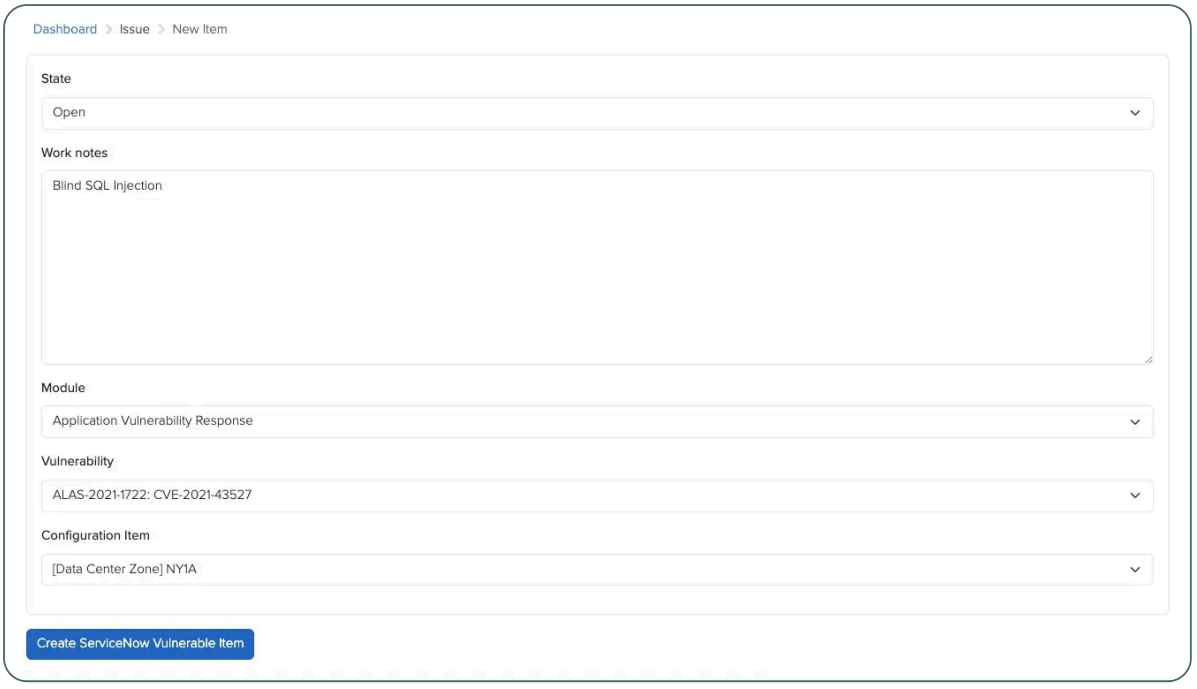
Dradis

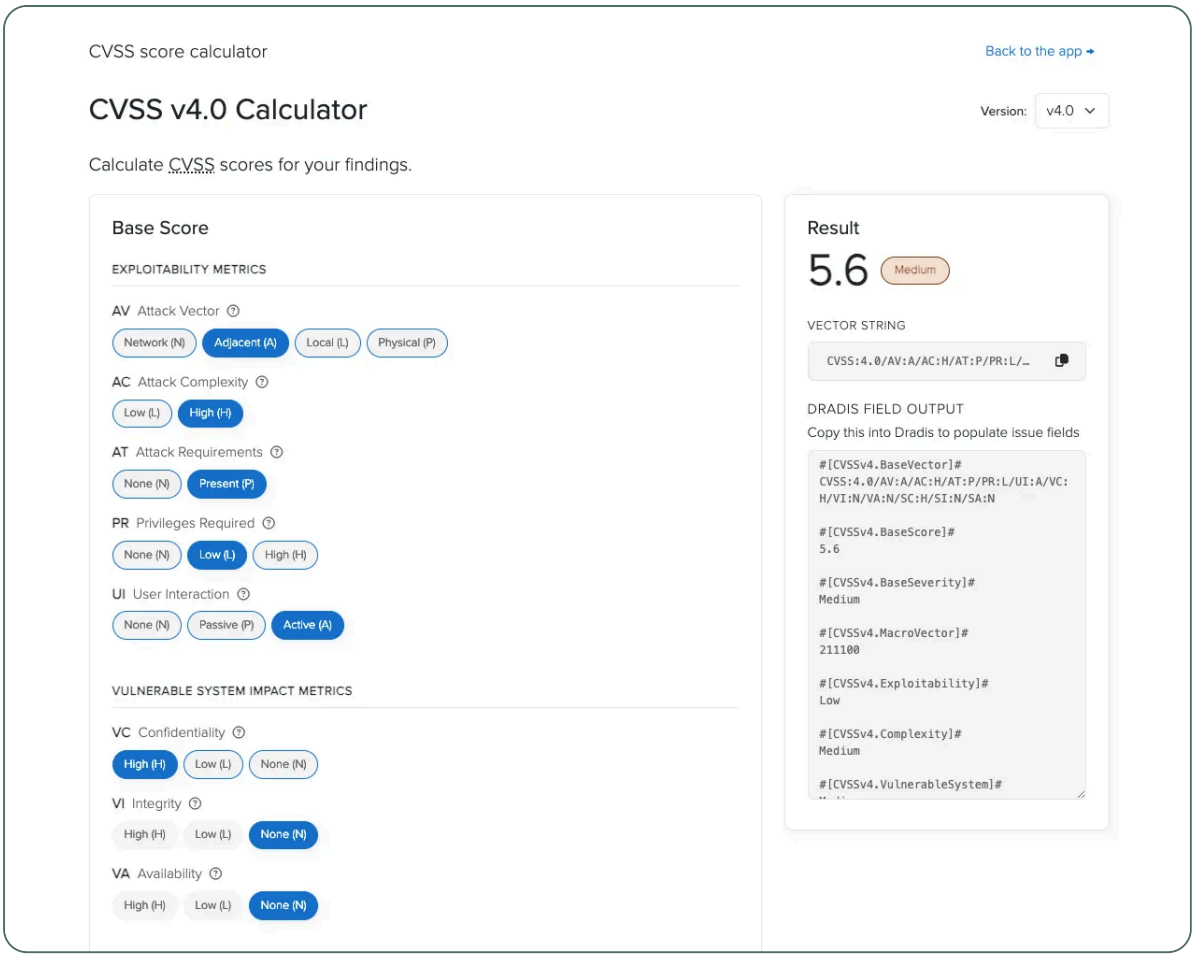
Feedback ⓘ New in Dradis Pro v5.2 (2026-07-09) BLOG ↗

Dradis is an open-source collaboration and reporting framework that helps security professionals consolidate findings, share information, and generate reports during penetration testing engagements.

Dradis Pro v5.2 adds List fields with multi-column issue sorting, broader ServiceNow module coverage, and a redesigned CVSS Calculator.







└──▷ **HOW TO FIND IT**

When you need to capture and paste a CVSS vector string into a ticket or report, use the redesigned calculator to score a finding and copy the vector in one click.

📍 In Dradis Pro, open a finding and click CVSS Calculator. Adjust metrics in the two-column layout — the score and severity badge update live in the sticky results panel. Click Copy vector string to copy the full CVSS vector to your clipboard.

New List field type lets you define ordered value sets (e.g., Critical > High > Medium > Low) for Issue fields, complete with configurable font and cell background colours.

Multi-column Issue sorting now lets you stack multiple List and numeric fields as ordered sort criteria, so reports emerge in exactly the priority sequence your methodology demands.

ServiceNow integration now creates Vulnerable Items across the Application, Container, and Vulnerability Response modules automatically — no extra configuration required.

Redesigned CVSS Calculator features a two-column layout with a sticky results panel, severity badges for instant visual context, and a one-click vector string copy button.

New default Welcome kit option added for Kits.

+2 more

WAS THIS USEFUL? 👍 👎

Bugcrowd

Feedback ⓘ VRT Update (1.19.1) (2026-07-08) BLOG ↗

Bugcrowd Customer Changelog is a feed that documents updates, new features, and improvements made to Bugcrowd's crowdsourced security platform for customers.

Bugcrowd VRT 1.19.1 adds Active Directory, Kerberos/SCCM, and server misconfiguration categories plus granular SSRF classifications.

Adds Active Directory misconfiguration entries to the VRT, covering a new class of findings for AD-related vulnerabilities.

Adds Kerberos and SCCM abuse vectors as distinct VRT categories, enabling more precise classification of common internal attack paths.

Adds server security misconfiguration classifications rated across the full P1–P5 severity spectrum.

Replaces legacy SSRF categories with granular SSRF metrics, enabling finer-grained severity differentiation for server-side request forgery findings.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Snyk

Feedback ⓘ Announcing Snyk CLI v1.1306.0 (2026-07-09)

5 RELEASES · 2026-06-16 → 2026-07-09 BLOG ↗ STABLE

Snyk Product Updates is an RSS feed that delivers the latest feature releases, improvements, and security updates across Snyk's developer security platform.

Snyk CLI v1.1306.0 adds a new `snyk doctor` diagnostic command, broader Java/JVM and .NET container scanning, and MCP breakability evaluation on by default.

└──▷ TRY IT

Quickly diagnose CLI configuration or environment issues before opening a support ticket.

```
$ snyk doctor
```

New `snyk doctor` command diagnoses common CLI problems by generating a diagnostic report for your system or analyzing debug log output.

Container scans now detect the Java runtime version across a wider range of JVM base images.

Container scans can now find vulnerabilities in .NET application dependencies.

The breakability evaluation tool in the Snyk MCP Server is now enabled by default and no longer requires an experimental flag.

Shows a warning when a request is automatically retried due to rate limiting, improving visibility into throttled scans.

▾ 4 more releases in this issue · 2026-06-16 → 2026-07-09

> Snyk Code: July Release, C++ rules, Java library coverage, and JavaScript Insecure Tra
>
> BLOG ↗ STABLE
>
> Snyk Code adds 6 new C++ rules, broader C++ native detection, expanded Java library coverage, and a JS/TS Insecure Transmission rule.
>
> New C++ rule: Log Forging (CWE-117, high) — detects untrusted user input reaching logging sinks, enabling log forgery or corruption.
>
> New C++ rule: Improper Privilege Management (CWE-269, high) — flags unverified privilege-dropping calls that may leave processes running with elevated privileges.
>
> New C++ rule: Missing Authorization (CWE-862/CWE-732) — detects world-writable/world-executable file permissions and root UID/GID passed to privilege-escalation or file-ownership functions.

New C++ rule: SSL/TLS Certificate Verification Bypass (CWE-295, medium) — detects disabled certificate verification across OpenSSL, Qt, mbedTLS, libcurl, Boost.Asio, libpq, and libpqxx.

New C++ rule: Insecure TLS Configuration (CWE-327, high) — detects enablement of outdated TLS versions and other insecure TLS settings.

+6 more

Snyk Learn lesson roundup: what's new in June (2026-06-17) BLOG ↗ STABLE

Snyk Learn adds three new AI security lessons on agentic RCE, memory poisoning, and inter-agent communication.

New AI-Sec lesson: Unexpected Code Execution (RCE) — covers how agentic systems can be tricked into executing unauthorized code.

New AI-Sec lesson: Memory and Context Poisoning — explains how corrupted agent memory can silently alter AI reasoning and decisions.

New AI-Sec lesson: Insecure Inter-Agent Communication — teaches exploitation of weak communication protections between agents.

Expanded OWASP Top 10 lesson coverage now available in Python, Rust, and Ruby.

Automatically Close Obsolete Open Source Fix PRs with Help from Snyk, Now Generally Av

BLOG ↗ STABLE

Snyk now auto-closes stale Open Source Fix PRs by default, with a configurable daily closure limit.

Enables automatic closing of obsolete Open Source Fix PRs for all organizations by default, triggered during recurring tests when the targeted dependency is removed, transitively updated, or manually fixed.

Adds a configurable maximum number of obsolete PRs Snyk will close per day, giving teams rate control over automated closures.

Allows administrators to opt out of automatic PR closure via the organization settings page.

Support for Bruno Collections in API Targets BLOG ↗ LATEST STABLE

Snyk API & Web now supports Bruno collections natively, enabling direct import to create API Targets.

└──▷ HOW TO FIND IT

Import a Bruno collection to create an API Target without pre-converting it to Postman or OpenAPI format.

> 📍 In the Snyk console, go to API & Web › API Targets, select Import Collection, and choose your Bruno collection file.

WAS THIS USEFUL?  

StackHawk

Feedback ⓘ snapshot-20260713 NOTES ↗

StackHawk is a dynamic application security testing (DAST) tool that helps developers find and fix security vulnerabilities in their web applications and APIs during development.

StackHawk adds standalone native binary, script validation daemon, org-wide findings table, service keys, and FlightPath AI auth config.

└──▷ **TRY IT**

Iterate on an auth script locally without triggering a full scan — daemon mode re-runs the script every time you save.

```
$ hawk validate script --watch --app-id <your-app-id> --target-url
https://app.example.com
```

Browse all available HawkScan config sections and then inspect the exact schema for the authentication block before writing your stackhawk.yml.

```
$ hawk config list && hawk config show app.authentication
```

Provision a CI service key scoped to Scan Only and store it for pipeline use — without tying it to any individual user account.

> 📍 In the StackHawk console, go to Settings › Service Keys and click Create Service Key. Copy the generated key into your CI secret store (e.g., as `HAWK_API_KEY`). The key belongs to the org and survives team member changes.

Adds `hawk validate script` command to test active and passive HawkScripts against a live URL without running a full scan; pass `--watch` for daemon mode that re-executes on every file save.

Adds standalone native HawkScan binary (macOS, Linux, Windows) with bundled JRE — no Java installation required in CI or on developer machines.

Integrates `hawk op` directly into the `hawk` binary, eliminating the need for a separately installed `hawkop` binary for platform operations.

Adds OS-native certificate trust (macOS Keychain, Windows Certificate Store, Linux PEM bundle) for zero-config support of transparent TLS-intercepting proxies (Zscaler, Netskope, corporate CAs); disable with `--no-native-truststore`.

Adds `--browser` flag to `hawk init` for non-interactive device-flow authentication — opens the browser immediately and prints the auth URL to stdout, designed for AI coding agents.

+11 more

└──▷ **BREAKING ON UPGRADE**

! The `hawk-{version}.zip` download format is no longer available; pipelines or scripts that download the ZIP directly must migrate to the OS- and architecture-specific binary download at `https://download.stackhawk.com/hawk/{VERSION}/{PLATFORM}/hawk` (Windows: `hawk.exe`).

! The StackHawk CircleCI Orb v2.0.0 renames jobs and parameters from kebab-case to snake_case (e.g., `stackhawk/hawkscan-local` → `stackhawk/hawkscan_local`, `app-id` → `app_id`); existing pipelines must update job and parameter names when upgrading to v2.0.0.

WAS THIS USEFUL? 👍 👎

Promptfoo

Feedback ⓘ snapshot-20260713 NOTES ↗

Promptfoo is an open-source tool for testing, evaluating, and red-teaming LLM applications, with release notes documenting its latest features and fixes.

Promptfoo January 2026 snapshot adds adaptive rate limiting, Transformers.js local inference, telecom red team plugins, and multi-input red team scans.

└──▷ **TRY IT**

Run a red team scan against a system with multiple input variables to cover complex prompt structures.

```
$ promptfoo redteam scan --env-file .env.prod --env-file .env.secrets
```

Add a description to a generated red team test suite for traceability in CI.

```
$ promptfoo redteam generate -d "Q3 telecom safety audit" --plugins
telecom
```

View recent Promptfoo log output to diagnose provider or scan issues.

```
$ promptfoo logs
```

Adds Transformers.js provider to run Hugging Face models locally in Node.js or the browser.

Adds Vercel AI Gateway and Cloudflare AI Gateway providers for request routing.

Adds AWS Bedrock Video (Nova Reel, Luma Ray 2) and Azure AI Foundry Video (Sora) providers for video generation.

Adds native session endpoint support to the HTTP provider for stateful multi-turn conversations.

Adds apiBaseUrl, websocketUrl override, and function call support to the xAI Voice provider.

+22 more

WAS THIS USEFUL? 👍 👎

Portkey

Feedback ⓘ snapshot-20260713 NOTES ↗

Portkey Changelog is a documentation page that tracks and communicates updates, new features, bug fixes, and improvements made to the Portkey AI gateway platform.

Portkey adds vault-backed credentials, weekly rate limits, cross-cloud log delivery, and four new guardrail integrations in March 2026.

└──▷ **USE IT**

Keep LLM provider keys out of Portkey by pointing a virtual key at a secret stored in AWS Secrets Manager — Portkey fetches the value at request time.

json

```
{
  "virtual_key": "my-openai-key",
  "secret_mappings": {
    "api_key": "arn:aws:secretsmanager:us-east-
1:123456789012:secret:openai-api-key-AbCdEf"
  }
}
```

Gate every request so it must carry a `user_id` metadata field — the gateway blocks the call before it hits the model if the field is absent.

json

```
{
  "guardrails": [
    {
      "type": "required_metadata",
      "required_keys": ["user_id", "session_id"]
    }
  ]
}
```

Apply your existing Zscaler Detection Policy to LLM inputs and outputs by wiring AI Guard into the gateway's request/response hooks.

```json
{
  "guardrails": [
    {
      "type": "zscaler",
      "hooks": ["beforeRequestHook", "afterRequestHook"],
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    }
  ]
}
```

Adds Secret References so Portkey fetches credentials at runtime from AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault via `secret_mappings`, keeping keys out of Portkey directly.

Supports weekly budget/usage windows (`rpw`) and endpoint-scoped rate limits, letting teams cap chat completions, embeddings, and admin routes independently.

Enables GCS log delivery from AWS-hosted gateways via GCP Workload Identity Federation (`GCP_WIF_AUDIENCE` / `GCP_WIF_SERVICE_ACCOUNT_EMAIL`), eliminating long-lived cross-cloud keys.

Allows organization admins to include archived workspaces in analytics graphs and filters, with workspace slugs accepted alongside IDs in saved filters.

Integrates Zscaler AI Guard as a guardrail partner, applying Zscaler Detection Policies to LLM inputs and outputs via `beforeRequestHook` / `afterRequestHook` with a required `policyId`.

+8 more

WAS THIS USEFUL? 👍 👎

Openlayer

Feedback ⓘ Batch test re-runs and expanded integrations

15 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Openlayer is an AI testing and monitoring platform that helps teams evaluate, debug, and improve machine learning models and LLM applications in production.

Openlayer adds batch test re-runs, Ruby SDK, Gemini/Copilot Studio integrations, and new drift detection.

.

└──▷ **TRY IT**

Retrieve test results programmatically to pipe into downstream alerting or reporting pipelines.

```
$ curl -X GET 'https://api.openlayer.com/v1/projects/<project-id>/test-results' \
  -H 'Authorization: Bearer <api-key>'
```

Adds batch and manual test re-runs so results can be recalculated retroactively when new data is added.

Supports automatic threshold configuration for tests built on custom metrics.

Introduces custom time range selection for filtering test execution history.

New primary key uniqueness constraint test type available on the platform.

Adds Ruby SDK for instrumenting Openlayer in Ruby applications.

+11 more

▾ **14 more releases in this issue** · seen 2026-07-13

> Complete design system overhaul, Snowflake Integration NOTES ↗ STABLE
>
> Openlayer adds Snowflake integration, new security role, prompt injection tests, and expanded LLM tracing for Bedrock, OpenAI Agents, Pydantic AI, and LangGraph.
>
> .
>
> └──▷ **HOW TO FIND IT**
>
> Connect a Snowflake table directly to a project to run evals against live warehouse data without manual exports.
>
> > 📍 In the console, open your project, go to Data Sources › Connect, select Snowflake, and choose the target table.
>
> Enable a prompt injection test on an LLM data source to automatically detect adversarial inputs in production traffic.

📍 In the console, open your project's Tests tab, click Add Test, and select `Prompt Injection` from the Evals library.

New UI with full light and dark mode support, updated design system, and improved navigation hierarchy.

Adds priority levels for tests and more data visualizations throughout the product.

SDK support for tracing Bedrock models.

SDK support for tracing OpenAI Agents.

SDK support for tracing Pydantic AI systems.

`+12 more`

└──▷ **BREAKING ON UPGRADE**

! `Inference pipelines` are renamed to `data sources` throughout the platform

`Custom metrics, rotating API keys, and new models for direct-to-API calls`

NOTES ↗ **STABLE**

Openlayer adds custom metrics via openlayer.json, rotating API keys, and 11 new models for direct-to-API evaluation calls.

.

└──▷ **USE IT**

Define a custom domain-specific metric in openlayer.json so it appears as a project-wide metric and a runnable test on all your data.

json

```
{
  "metrics": [
    {
      "name": "toxicity_score",
      "description": "Custom toxicity scorer for our domain",
      "type": "custom",
      "script": "metrics/toxicity.py"
    }
  ]
}
```

Rotate API keys by creating a new key before revoking the old one — navigate to your account API settings to manage multiple keys.

📍 In the Openlayer console, go to Account Settings › API Keys, click `Create new key`, copy the new key into your services, then delete the old key.

Supports custom metrics uploaded via `openlayer.json`, usable as project-wide metrics or as runnable eval tests alongside standard metrics.

Enables creation, renaming, and deletion of multiple personal API keys to support API key rotation workflows.

Adds ability to specify desired metrics directly in `openlayer.json` for API-driven evaluation runs.

Adds 11 new models for direct-to-API calls: GPT-4o, GPT-4 Turbo, Claude 3.5 Sonnet, Claude 3 Haiku, Claude 3 Opus, Claude 3 Sonnet, Command R, Command R Plus, Gemini 1.0 Flash, Gemini 1.5 Flash, and Gemini 1.5 Pro.

Adds git repository and organization linking in git settings pages for tighter version-control integration.

+1 more

Improved quality control over your LLM's responses with annotations and human feedback

NOTES ↗ STABLE

Openlayer adds request annotations, human feedback ratings, data export, and Git integration to LLM quality control.

.

└──▷ **HOW TO FIND IT**

Review a filtered set of low-quality responses offline by exporting only the flagged requests cohort from your workspace.

> 📍 In the console, go to your workspace › Requests, apply the desired filters (e.g., thumbs-down ratings), then click Export to download the filtered cohort as a file.

Connect a project to a Git repository and configure the target branch and root directory for version-controlled evaluation pipelines.

> 📍 In the console, go to your Project › Settings › Git Repository, enter the repository URL, set the branch and root directory, then save.

Adds per-request annotations with custom values, including step-level annotation for individual trace steps when tracing is enabled.

Adds thumbs up/thumbs down human feedback ratings on every request to quickly surface good and bad model responses.

Enables data export from the UI, including filtered cohorts of request data.

Introduces updated navigation with breadcrumbs and the ability to copy project name and inference pipeline ID.

Adds project-level Git repository settings, including editable branch and root

directory configuration.

+4 more

Improved test diagnosis page, SAML SSO, design refreshes, + more NOTES ↗

STABLE

Openlayer adds SAML SSO, Azure OpenAI evaluator, Ragas metrics, and an overhauled test diagnosis page.

.

└──▷ HOW TO FIND IT

Enforce org-level SSO for your team's Openlayer workspace to add a layer of authentication security.

📍 In the console, go to Settings › Collaboration › SAML SSO and configure your identity provider credentials.

Adds SAML SSO support via any major identity provider, enabling teams to enforce org-level authentication for workspace access.

Supports Azure OpenAI as an LLM evaluator integration.

Adds additional Ragas metrics — faithfulness and answer correctness — for evaluation.

New list view on the test diagnosis page surfaces error messages for skipped and errored tests, with filtering by result type and a status-breakdown overview.

Test diagnosis page sections are now collapsible for a streamlined review experience.

+4 more

Introducing Openlayer Governance NOTES ↗ STABLE

Openlayer adds Governance mode, A/B testing, new data backends, CLI data sources, and expanded tracing support.

.

└──▷ HOW TO FIND IT

Enable Governance mode to define and track enforcement rules across your workspace's AI projects.

📍 In the Openlayer console, open your workspace, navigate to Governance, select New Framework, define rules (e.g. PII leakage test, monitoring required, risk level), assign owners, and track compliance across projects.

New Governance mode lets teams define and enforce custom policy frameworks (PII leakage, prompt injection, monitoring requirements, ownership, risk level) across all projects and workspaces, with rule owners, comment threads, and evidence uploads.

New workspace overview and project inventory gives a birds-eye view of project activity with categorization and labeling across the workspace.

Adds support for Databricks, Redshift, and Postgres as third-party data backends in Monitoring mode.

Enables A/B testing across two commits (Development mode) or two data sources (Monitoring mode) to compare performance across metrics.

Supports configuring directory sync for SAML/SSO-enabled workspaces via the UI.

+12 more

Openlayer Gateway, customizable dashboards, and Salesforce Agentforce GA

NOTES ↗ STABLE

Openlayer adds a managed AI gateway, customizable dashboards, Salesforce Agentforce GA, Helm on-prem, new guardrails, governance frameworks, and audit trails.

⋮

└──▷ **HOW TO FIND IT**

Route all LLM calls through the Openlayer Gateway to get centralized API key management, rate limiting, and end-to-end observability without changing your model provider code.

📍 In the console, go to Gateway › Admin Portal, create a gateway API key, assign rate limits and team, then point your LLM client's base URL to the Openlayer Gateway endpoint instead of the provider directly.

Delete stale monitoring records programmatically to keep your monitoring dataset clean.

```
$ curl -X DELETE https://api.openlayer.com/v1/monitoring/records \
  -H 'Authorization: Bearer <your-api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"recordIds": ["<record-id-1>", "<record-id-2>"]}'
```

Deploy Openlayer on-prem to your Kubernetes cluster using the new Helm chart from Docker Hub OCI.

```
$ helm install openlayer oci://registry-
1.docker.io/openlayer/openlayer \
  --set storage.type=s3 \
  --set database.url=<your-db-url> \
  --set replicaCount=3 \
  --set image.pullSecrets[0]=<pull-secret> \
  --set node.heapSize=4096
```

New Openlayer Gateway: centralized managed control layer for all LLM calls, with admin portal, API key management, per-key rate limits and freeze, and team-based organization.

Customizable project home pages: pin metric widgets, view project notifications inline, and navigate alerts from a new sidebar notification center.

Salesforce Agentforce integration reaches GA, with configurable login domain (production, sandbox, or My Domain) and updated Builder schema capturing every intermediate step.

On-prem Helm chart published to Docker Hub OCI for Kubernetes deployments, with config for storage, database, replicas, image pull secrets, and Node heap.

Built-in governance frameworks available out of the box: EU AI Act, TRAIGA, E23 MRM, and the Openlayer Governance Framework.

+30 more

Pausing tests, checks for duplicate keys, and new integrations NOTES ↗

STABLE

Openlayer adds Google ADK integration, Snowflake remote execution, pausing tests, duplicate-key checks, and cloud marketplace availability.

.

Adds column distribution graphs for LLM projects on the platform.

New SDK integration with Google Agent Development Kit (ADK) for evaluating multi-agent systems.

Enables pausing test execution in Monitoring mode.

New CLI command to export tests for a project.

New test to check for duplicate unique and primary keys.

+8 more

Project-level secrets, tracing LLM requests with OpenTelemetry NOTES ↗

STABLE

Openlayer adds project-level secrets, OpenTelemetry tracing, Azure OpenAI for evals, and new platform APIs.

.

└──▷ **HOW TO FIND IT**

Store a sensitive API key once at the project level so all evals and tests can reference it without hardcoding credentials.

> 📍 In the Openlayer console, go to your Project › Settings › Secrets and add a new secret (e.g., `OPENAI_API_KEY`). Reference it by name in your test configurations.

Set a deployment-wide default LLM-as-a-judge API key for all on-prem workspaces so teams don't need to configure keys individually.

> 📍 In the Openlayer on-prem admin console, go to Deployment Settings › Default API Keys and set the key for LLM-as-a-judge.

Adds project-level secrets storage in the Platform so API keys and auth tokens can be securely referenced across tests without hardcoding.

Enables setting default LLM-as-a-judge API keys for an entire on-premise deployment, eliminating per-workspace/project key configuration.

Adds SDK support for tracing LLM requests with OpenTelemetry.

Supports Azure OpenAI deployments as the backend for LLM-as-a-judge evaluations.

Adds gpt-4o-mini as an available LLM evaluator model.

+6 more

SAML Directory Sync, new LLM-as-a-judge models, and website refresh **NOTES ↗**

STABLE

Openlayer adds SAML directory sync, Cohere/Vertex AI LLM-as-a-judge models, new RAG metrics, Viewer role, and expanded APIs.

.

└──▷ **TRY IT**

Include project and workspace details in a single inference pipeline GET call, avoiding extra round-trips.

```
$ curl -X GET 'https://api.openlayer.com/v1/inference-
pipelines/<id>?expand=project,workspace' \
  -H 'Authorization: Bearer <token>'
```

Check how many tests are skipped or failing on a project version without fetching full test details.

```
$ curl -X GET 'https://api.openlayer.com/v1/projects/<project-id>/versions/<version-id>' \
  -H 'Authorization: Bearer <token>'
```

Enable SAML directory sync so workspace membership mirrors your identity provider's security groups.

> 📍 In the Openlayer console, go to Workspace Settings › Security › SAML SSO and enable Directory Sync to map your IdP groups to Openlayer workspaces.

Supports SAML SSO directory sync so Openlayer workspaces can be governed by existing organizational security groups.

Adds Cohere and Vertex AI models as judge options for LLM-as-a-judge evaluation tests.

New Faithfulness and Answer Correctness metrics available for RAG system evaluations.

New `Viewer` workspace role with read-only access — no write, update, or delete permissions on resources.

Adds `expand` parameter to inference pipeline GET endpoints, embedding project and workspace data directly in the response body.

+8 more

Stronger security, multimodal tracing, and expanded governance controls

NOTES ↗ STABLE

Openlayer adds MFA, multimodal traces, directory sync, MCP integration, and expanded API/CLI controls for AI governance.

.

└──▷ **TRY IT**

Force a metrics run in CI to ensure all metrics execute even when no changes are detected.

```
$ openlayer metrics run -force
```

Adds platform-wide multi-factor authentication (MFA) for enterprise security.

Supports deployment-wide directory sync with project-level access group support, including a NoAccess group level for finer-grained access control.

Renders images, audio, and PDFs inline in trace steps and data tables for richer multimodal observability.

Surfaces rule-result summaries on Governance and Frameworks pages to streamline audit readiness and compliance reporting.

New Langflow integration for connecting Langflow-based AI workflows.

+17 more

└─▷ **BREAKING ON UPGRADE**

! Openlayer system tags are renamed to `categories`; user-defined labels remain `tags`, which may break any workflows, scripts, or integrations that reference the old `tags` name for system tags

`Support for tracking users and sessions, test tags, and more` **NOTES ↗**

STABLE

Openlayer adds user/session tracking, test tags, LiteLLM support, PII redaction guardrails, and multi-agent LangGraph workflows.

.

└─▷ **HOW TO FIND IT**

Redact PII automatically from traces in production before they reach Openlayer.

📍 In the Openlayer SDK, enable the guardrail system for tracing by configuring the guardrails option to redact PII on your tracer initialization.

Control which AWS Signature version Openlayer uses in an on-prem deployment without changing application code.

```
$ export AWS_SIGNATURE_VERSION=v4
```

Set up user and session tracking to view full trace history per user in the Openlayer app.

📍 In the Openlayer app, go to the Data section and enable user and session tracking to associate traces with individual users and sessions.

Adds user and session tracking in the Data section, enabling full session and record history per user for behavioral analysis.

Adds predefined and custom tags on tests to browse and organize tests by outcome, domain (e.g., safety, data drift), or governance framework (e.g., EU AI Act).

New API support for user-defined tags on tests and filtering tests by those tags.

New API support for batch LLM calls.

New API support for user sessions to track traces across interactions.

+8 more

Test bundles, new tests, support for new Python runtimes NOTES ↗ STABLE

Openlayer adds test bundles (EU AI Act, OWASP, agentic), new LLM tests, and expanded Python runtime support.

.

└──▷ HOW TO FIND IT

Apply a pre-built EU AI Act test bundle to catch compliance gaps across your LLM pipeline before deploying.

📍 In the Openlayer app, go to Tests › Test Bundles, select the EU AI Act bundle, and click Apply to your project.

Filter your test results to surface only high-priority failing tests during an incident triage.

📍 In the Openlayer app, go to Tests, then use the Filter controls to set Priority = High and Status = Failing.

Introduces test bundles grouping related tests by policy or use case — pre-built sets for EU AI Act, OWASP, agentic workflows, and data quality.

Adds new tests to the library: groundedness for LLMs, toxicity score, prompt injection detection, and competitor-recommendation checks.

Supports new Python runtimes for CLI development mode.

Adds filtering of tests by priority, status, name, and other attributes in the platform UI.

Updates OpenTelemetry GenAI semantic conventions attribute handling in the API.

+2 more

The Openlayer MCP server, Automatic thresholds, BigQuery Integration and Anomaly Detec

NOTES ↗ STABLE

Openlayer adds MCP server, automatic thresholds, BigQuery integration, anomaly detection, and project-level access groups.

.

└──▷ HOW TO FIND IT

Let Openlayer automatically determine acceptable thresholds for a new test instead of guessing manually — ideal when you lack historical baselines.

📍 In the Openlayer app, go to your project › Tests › Create new test. In the threshold configuration step, select `automatic` instead of entering a manual value. Openlayer

will derive thresholds from historical data patterns for that metric.

Connect a BigQuery table to a project for ongoing, automated data quality monitoring without manual data uploads.

📍 In the Openlayer app, go to your project › Integrations › BigQuery and follow the prompts to link a BigQuery table. Once connected, Openlayer will continuously pull data and run data quality tests against it.

Run Openlayer tests and explore results without leaving your IDE by using the MCP server.

📍 In your IDE's MCP configuration, add the Openlayer MCP server as a provider. Once connected, invoke Openlayer tests directly from the IDE workflow to get evaluation results inline during development.

New Openlayer MCP server enables use of Openlayer tests directly inside IDE workflows.

Supports automatic, data-driven thresholds for tests that adapt over time based on historical patterns — covering cost, data quality, and GPT eval scores.

Adds anomaly detection to flag test results that deviate from predicted norms.

New column statistic test supports mean, median, min, max, std, sum, count, and variance.

New raw SQL query test for custom evaluation logic.

+21 more

Trace timeline, project lifecycles, and Claude Agent SDK tracing NOTES ↗

STABLE

Openlayer adds Gantt trace timelines, Claude Agent SDK tracing, auto-instrumentation, and Gateway guardrails enforcement.

.
.

└──▷ **USE IT**

Bootstrap full automatic LLM SDK instrumentation for all supported providers in one call, then run your app — no per-SDK setup needed.

```python
import openlayer
openlayer.init(api_key="<your-api-key>", project_name="<project-name>")
# All installed supported SDKs (OpenAI, Anthropic, Mistral, Groq, Gemini, LiteLLM, Portkey, Google ADK, Azure) are now patched automatically.
```

Opt out of automatic instrumentation when you want manual control over which SDKs are traced.

```python
import openlayer
openlayer.init(api_key="<your-api-key>", project_name="<project-name>", auto_instrument=False)
```

Inspect the full agent loop — tool calls, MCP tools, subagent nesting, and token costs — by enabling Claude Agent SDK tracing in Python.

```python
import openlayer
import anthropic

openlayer.init(api_key="<your-api-key>", project_name="<project-name>")
# The Claude Agent SDK is now instrumented: tool calls, MCP tools, subagents,
# session continuity, and cost/tokens will appear in Openlayer traces.
```

New Gantt-style timeline view for traces — every span's start time, duration, and nesting shown side by side, making slow steps and unexpected branches immediately visible.

Adds project lifecycle stages and evidence document attachments for governance tracking, living alongside tests and results on each project.

First-party tracing for the Anthropic Claude Agent SDK in Python — captures tool calls, MCP tools, subagent nesting, session continuity, and cost and token counts per run.

New automatic SDK instrumentation via openlayer.init() — one call patches every installed supported LLM SDK (OpenAI, Anthropic, Mistral, Groq, Gemini, LiteLLM, Portkey, Google ADK, Azure); opt out with `auto_instrument=False`.

Gateway can now enforce guardrails policies inline on LLM requests routing through it.

+9 more

└──▷ **BREAKING ON UPGRADE**

- ! The configure() function is deprecated in favor of init(); trace configuration is now unified around openlayer.init(), with settings merging incrementally instead of resetting on each call.
- ! The Copilot Studio Logic App and Batch API integration paths are now deprecated in favor of the pull-based native data source.

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ Dev Machine Guard v1.13.0 (2026-07-09)

7 RELEASES · 2026-06-16 → 2026-07-09 BLOG ↗ LATEST STABLE

StepSecurity Changelog is an RSS feed that tracks and publishes updates, new features, and improvements to the StepSecurity platform for securing CI/CD pipelines.

Dev Machine Guard v1.13.0 makes Device Policy enforcement GA and adds Visual Studio discovery and disk-based package scanning.

└──▷ HOW TO FIND IT

Assign a Device Policy profile to enforce an approved VS Code extension set across all developer machines without MDM.

> 📍 In the dashboard, go to Device Policy › Profiles, create or select a profile, then assign it to All Devices. The agent installs and continuously enforces the approved extension set on its next cycle.

Device Policy enforcement via the DMG agent is now generally available — enforces approved VS Code extension sets on developer machines without requiring MDM, assignable per device or globally from the dashboard.

New classic Visual Studio installation discovery extends IDE extension inventory beyond VS Code to legacy Visual Studio versions and their extensions.

New disk-based package scanning for npm (lockfile parsing) and Python (dist-info metadata) with a scan cache and concurrent workers for faster runs; previous method remains available via the `use_legacy_package_scan` opt-out.

Scans can now be triggered at login, with improved scheduling reliability and scheduler diagnostics.

Expanded global package discovery adds support for pnpm v10+ globals and Yarn classic globals on Windows.

▾ 6 more releases in this issue · 2026-06-16 → 2026-07-09

> Harden Runner Releases- July 7, 2026 (2026-07-07) BLOG ↗ STABLE
>
> Harden Runner adds egress block mode for macOS/Windows, Bitrise support, and Bun HTTPS monitoring.
>
> └──▷ USE IT
>
> Restrict outbound traffic on a Windows GitHub-hosted runner to only allowed endpoints — now that block mode is supported beyond Linux.

```yaml
- uses: step-security/harden-runner@v2.20.0
  with:
    egress-policy: block
    allowed-endpoints: >
      api.nuget.org:443
      github.com:443
```

Supports `egress-policy: block` mode on macOS and Windows GitHub-hosted runners, enabling outbound traffic restriction (not just auditing) on these platforms.

Adds compatibility with Bitrise-hosted macOS runners for GitHub Actions workflows.

Enables HTTPS request monitoring for workflows using the Bun runtime on Linux runners (enterprise tier) in the GitHub Action.

Enables HTTPS request monitoring for Bun on Linux amd64 self-hosted VM runners.

New Detection: Secret Exfiltration Attempt (2026-06-30) BLOG ↗ STABLE

Harden-Runner adds Secret Exfiltration Attempt detection to catch toJSON(secrets) supply-chain attacks in real time.

└──▷ **HOW TO FIND IT**

Get an immediate alert when a secret exfiltration attempt is detected, so you can rotate exposed tokens before an attacker retrieves the artifact.

> 📍 In the Harden-Runner console, go to Notification Settings and enable the `Secret Exfiltration Attempt` detection notification.

New `Secret Exfiltration Attempt` detection type flags workflow runs that combine known attack indicators — malicious workflow name, toJSON(secrets) secret dump, suspicious commit message, and an exfiltration artifact — before stolen secrets can be retrieved.

Confidence scoring elevates a detection to High when the workflow run, branch, or workflow file is deleted shortly after execution, treating evidence removal as a strong exploitation signal.

Direct navigation in the Detections page pinpoints the affected repository, matched rule, indicators, workflow file, and the exact job, step, and line where the secret dump occurred.

Opt-in alerting available in Notification Settings to receive an immediate notification the moment a Secret Exfiltration Attempt is flagged.

StepSecurity Maintained Actions Are Now Free for Public Repositories (2026-06-29)

BLOG ↗ STABLE

StepSecurity Maintained Actions — secure drop-in replacements for 500+ GitHub Actions — are now free for public repositories.

Harden Runner Releases- June 24, 2026 (2026-06-24) BLOG ↗ STABLE

ARC Harden Runner v2.18.3 adds broker pod customization fields and Cosign keyless image signing.

└──▷ TRY IT

Verify the arc-harden-runner container image's provenance before deploying, ensuring it was built by the official GitHub Actions pipeline.

```
$ cosign verify ghcr.io/step-security/arc-harden-runner/arc-harden-runner:v2.18.3 \
  --certificate-identity-regexp 'https://github.com/step-security/arc-harden-runner/.*' \
  --certificate-oidc-issuer https://token.actions.githubusercontent.com
```

Verify the broker sidecar image's provenance to confirm supply-chain integrity before cluster deployment.

```
$ cosign verify ghcr.io/step-security/broker/broker:v1.3.0 \
  --certificate-identity-regexp 'https://github.com/step-security/broker/.*' \
  --certificate-oidc-issuer https://token.actions.githubusercontent.com
```

Adds broker pod-spec customization via new Helm values under `broker.*` (volumeMounts, imagePullSecrets, automountServiceAccountToken, serviceAccountName, resources, nodeSelector, tolerations, affinity, topologySpreadConstraints) for scheduling alongside strict workloads.

Supports Cosign keyless (OIDC) signing for versioned and FIPS arc-harden-runner and broker image manifests published to ghcr.io, enabling supply-chain verification.

Secure Registry: Source Attribution, Typosquatting Protection, and Sonatype Nexus Supp

BLOG ↗ STABLE

Secure Registry gains source attribution, typosquatting blocking, and Sonatype Nexus Repository support.

└──▷ **HOW TO FIND IT**

Attribute CI pipeline package requests to a specific GitHub Actions run so Policy Evaluations log entries link back to the correct workflow.

> 📍 Set your Secure Registry API key token to the form <token>::gha:<identifier> in your pipeline's package manager configuration, following the Setup Guide.

Route npm traffic through Sonatype Nexus Repository while enforcing Secure Registry policy, using the new Nexus integration path.

> 📍 In the Secure Registry Setup Guide, select Sonatype Nexus Repository as the integration path and configure Nexus to proxy npm through your Secure Registry endpoint.

Enable Typosquatting Protection for an ecosystem and allowlist a legitimate package that triggers a false positive.

> 📍 In Secure Registry settings, enable Typosquatting Protection for the target ecosystem, then add the exact package name to the Whitelisted Packages allowlist to permit it.

Adds Source Attribution to the Policy Evaluations log via two new columns — Source and Source Identifier — tracing every package request to the developer machine or CI pipeline that made it.

Enables attribution opt-in by appending an identifier suffix to your API key token (`::dev:` for developer machines, `gha:` for GitHub Actions pipeline runs).

New Typosquatting Protection control blocks packages whose names closely resemble popular packages, configurable per ecosystem with a Whitelisted Packages allowlist to clear false positives by exact package name.

Adds Sonatype Nexus Repository as an integration path for teams that proxy npm through Nexus, joining JFrog Artifactory, Google Artifact Registry, and Direct npm.

`Maven Ecosystem Support in GitHub Checks and OSS Package Search (2026-06-16)`

`BLOG ↗` `STABLE`

StepSecurity adds Maven (Java) ecosystem support to GitHub Checks PR controls and OSS Package Search.

└──▷ **HOW TO FIND IT**

Enable both new Maven PR controls to protect your Java repositories from compromised or freshly published dependencies.

> 📍 In the StepSecurity console, go to GitHub Checks › PR Content Security, select your Maven repository, and enable Maven Package Compromised Updates and Maven

Package Cooldown — setting each to Required or Optional as your policy requires.

Find where a specific Maven package appears across PRs, default branches, and developer machines to assess blast radius before incident response.

📍 In the StepSecurity console, go to OSS Package Search, select Maven from the Package ecosystem toggle, enter the package coordinates (e.g. `com.fasterxml.jackson.core`:jackson-databind), and review results across pull requests, default branches, and developer machines.

Adds Maven Package Compromised Updates PR control: blocks pull requests that introduce or update a Maven dependency flagged in StepSecurity's compromised-packages database.

Adds Maven Package Cooldown PR control: blocks pull requests adding recently published Maven versions until they age past the configured cooldown window.

Enables Maven ecosystem selection in OSS Package Search, allowing blast-radius analysis of specific Maven packages across pull requests, default branches, and developer machines.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Training updates (2026-07-08)

5 RELEASES · 2026-06-24 → 2026-07-08 BLOG ↗ LATEST STABLE

Veracode Docs Updates is an RSS feed that delivers real-time notifications about the latest changes and additions to Veracode's official documentation.

Veracode Security Labs adds OWASP LLM 07 lesson and expands eLearning captions for C/C++ memory management course.

New Security Labs lesson: OWASP LLM 07: Under the Hood (Flask), expanding the OWASP Top 10 for LLM Applications hands-on lab series.

Adds multilingual caption support (English, Spanish, French, Portuguese, Chinese) to the C/C++ Memory Management Risks and Best Practices eLearning course.

▾ 4 more releases in this issue · 2026-06-24 → 2026-07-08

Integration updates (2026-07-07)

BLOG ↗ STABLE

GitHub Workflow Integration v2.0.0 adds platform-based install, scan scheduling, and Apex/VB6/Perl autopackaging support.

└──▷ USE IT

Enable debug logging in a GitHub or Azure DevOps workflow to troubleshoot scan failures.

```yaml
# In veracode.yml
debug: true
```

Target a specific analysis branch for scheduled scans in Azure DevOps instead of the default branch.

```yaml
# In veracode.yml
analysis_branch: main-security-baseline
```

Install and configure the GitHub Workflow Integration directly from the Veracode Platform to skip manual GitHub Marketplace steps.

> 📍 In the Veracode Platform, go to Integrations › GitHub › Install, select your GitHub organization, and complete the configuration to install or update credentials without visiting GitHub Marketplace.

Enables installing and configuring the Veracode GitHub Integration directly from the Veracode Platform, eliminating manual GitHub Marketplace steps.

Supports updating integration credentials and adding organizations from the Veracode Platform, regardless of original install method.

Adds autopackaging support for Apex, VB6, and Perl in the GitHub Workflow Integration.

Enables scheduling Static Analysis, SCA, and IaC scans for GitHub organizations from the Veracode Platform.

Adds a Scan Scope setting in Veracode Scan for Visual Studio v0.8.0 to scan the full project directory or only the active solution.

+13 more

Platform updates (2026-07-06) BLOG ↗ STABLE

Veracode adds six new built-in policies, OWASP/CWE 2025 auto-updates, OAuth 2.0 credentials UI, CMEK, and expanded Reporting API filters.

└──▷ **TRY IT**

Verify authenticated connectivity to Veracode after migrating away from the removed `/healthcheck` /status endpoint.

```
$ curl -s 'https://api.veracode.com/api/authn/v2/principal' \
  -H 'Authorization: Bearer <token>'
```

Updates Auto-Update OWASP and Auto-Update CWE Top 25 policy rules to enforce OWASP Top 10:2025 and MITRE CWE Top 25 for 2025 standards.

Adds six new built-in security policies to the platform, ahead of deprecation of recommended and transitional policies.

Adds `collections_list` field to Finding Reporting API responses, returning collection names associated with each application.

Adds Collections fields to the Applications explore in Veracode Analytics, with new Collections visualizations in the Policy Compliance Overview dashboard.

New `most_recent_scan` filter for the Reporting API Scans report — when set to `true`, returns only the most recent scan per asset.

+17 more

DAST updates (2026-06-26) BLOG ↗ STABLE

Veracode DAST adds Directory Restrictions, enhanced scheduling with breaking API changes, and open-port scanning with new detections.

└──▷ **HOW TO FIND IT**

If your scheduling integration uses the old DAST REST API endpoints, you must update your calls to the new scheduling endpoints to avoid breakage after the June 24 migration.

> 📍 Refer to the DAST API specification for the updated scheduling endpoint paths, then update any REST calls that previously targeted the old schedule management endpoints (e.g., replace deprecated schedule routes with the new ones documented in the DAST API spec).

Adds Directory Restrictions to DAST target configuration, giving practitioners control over scan scope by limiting which directories are crawled.

Introduces enhanced DAST Schedule configuration with advanced scheduling options and greater flexibility; all existing schedules are automatically migrated.

Breaking change: DAST REST API scheduling endpoints have changed — integrations managing schedules via the REST API must be updated to the new endpoints.

Enables open-port scanning for web applications in Dynamic Analysis and DAST Enterprise mode.

Adds new scan engine detections for HTTP headers, TLS, and fingerprinting in Dynamic Analysis and DAST Enterprise mode.

Static Analysis updates (2026-06-24) BLOG ↗ STABLE

Veracode SAST June 2026 engine adds Dart 3.12/Flutter 3.44 and Kotlin 2.4 support, plus new CWE detections across multiple languages.

Adds support for Dart 3.12 and Flutter 3.44 in static analysis.

Adds support for Kotlin 2.4.

Improves classification of first-party .NET assemblies to better distinguish application code from third-party dependencies.

Improves detection of vulnerabilities involving tainted data flows through public Apex methods.

Improves support for Apple Security framework APIs.

+4 more

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

Cooldown policies now support NuGet (2026-07-07)

10 RELEASES · 2026-06-15 → 2026-07-07 BLOG ↗ LATEST STABLE

Cloudsmith Changelog is a feed tracking updates, new features, and improvements to Cloudsmith's cloud-native package management and artifact repository platform.

Cloudsmith cooldown policies now extend to NuGet packages for supply-chain protection.

Supports cooldown policies for NuGet packages, blocking newly published versions from being consumed until a configured quarantine period has elapsed.

▾ 9 more releases in this issue · 2026-06-15 → 2026-07-07

> Cloudsmith API keys are now detectable by GitHub secret scanning (2026-07-03)
>
> BLOG ↗ STABLE
>
> Cloudsmith API keys are now automatically detected by GitHub secret scanning to prevent unauthorized use.
>
> └──▷ HOW TO FIND IT
>
> Enable GitHub secret scanning on your repository so exposed Cloudsmith API keys are flagged automatically before they can be exploited.
>
> > 📍 In the GitHub repository, go to Settings › Security › Code security and analysis, then enable Secret scanning.

> Speed up Python builds with PEP-658 support (2026-07-02) BLOG ↗ STABLE
>
> Cloudsmith now supports PEP-658, letting Python clients fetch package metadata without downloading full wheels.

> Cooldown policies now support Go (2026-07-01) BLOG ↗ STABLE
>
> Cloudsmith cooldown policies now extend supply-chain protection to Go packages.
>
> Supports cooldown policies for Go packages, quarantining newly published versions until a configured delay period has elapsed.

> Set up Docker Hub and DHI upstreams with Cloudsmith-managed authentication (2026-06-29
>
> BLOG ↗ STABLE
>
> Cloudsmith now manages Docker Hub and DHI upstream auth automatically, raising rate limits without user-supplied credentials.

└──▷ **HOW TO FIND IT**

Configure a Docker Hub upstream proxy in Cloudsmith without supplying your own Docker Hub credentials to get higher pull rate limits.

> 📍 In the Cloudsmith console, go to your repository › Upstreams › Add Upstream › Docker Hub, and complete setup without entering credentials — Cloudsmith supplies a managed token automatically.

Supports Cloudsmith-managed authentication for Docker Hub upstreams, eliminating the need to provide personal credentials.

Supports Cloudsmith-managed authentication for Docker Hardened Images (DHI) upstreams with the same credential-free setup.

Enables higher Docker Hub and DHI pull rate limits automatically via managed authenticated tokens versus anonymous access.

`Enrich packages with custom metadata (2026-06-16)` **BLOG ↗** **STABLE**

Cloudsmith now supports attaching custom key-value metadata to packages for CI, provenance, compliance, and cost attribution.

`Vulnerability detection now covers ecosystem-native OSV advisories (2026-06-16)` **BLOG ↗** **STABLE**

Cloudsmith vulnerability detection now matches OSV advisories using ecosystem-native version ranges for Debian, Alpine, PyPI, and more.

`Manage policies as code with Terraform (2026-06-16)` **BLOG ↗** **STABLE**

Cloudsmith's Terraform provider v0.0.75 adds full lifecycle management of policies and policy actions as code.

`More events now trigger policy evaluations (2026-06-15)` **BLOG ↗** **STABLE**

Cloudsmith adds two new events that automatically trigger policy evaluations across repositories.

Adds two new policy evaluation triggers so policies are enforced automatically on additional repository events without manual intervention.

`Create and manage connected repositories with Terraform (2026-06-15)` **BLOG ↗** **STABLE**

Cloudsmith Terraform provider gains a new resource to define and manage connected repositories as code.

└──▷ **USE IT**

Declare a connected repository in Terraform to manage upstream proxy or mirroring configurations alongside the rest of your Cloudsmith infra.

```hcl
resource "cloudsmith_connected_repository" "example" {
  # configuration fields here
}
```

New `cloudsmith_connected_repository` Terraform resource enables defining connected repository configurations alongside other Cloudsmith infrastructure as code.

WAS THIS USEFUL? 👍 👎

Socket

Feedback ⓘ

Packagist Package Pages Now Surface Force-Pushed Tags (2026-07-06)

5 RELEASES · 2026-06-16 → 2026-07-06 BLOG ↗ LATEST STABLE

Socket Changelog is a feed tracking updates, improvements, and new features released by Socket, a security tool that protects software supply chains from malicious packages.

Socket now surfaces full commit history for force-pushed Packagist tags, highlighting the current and prior commits.

Displays full commit history for force-pushed tags on Packagist package pages, with the currently associated commit highlighted.

Enables review of commits a package version previously referenced before a tag was overwritten.

▾ 4 more releases in this issue · 2026-06-16 → 2026-07-06

New: AI-Reviewed Reachability Specifications (2026-06-30) BLOG ↗ STABLE

Socket adds opt-in AI-generated reachability specs to expand vulnerability reachability coverage beyond human-verified rules.

.

└──▷ HOW TO FIND IT

Enable AI-generated reachability specs org-wide to get reachability verdicts on a broader set of vulnerabilities during alert scans.

> 📍 In the console, go to Settings → Alert Scans → Reachability and enable the AI-generated reachability specifications option.

Adds opt-in AI-reviewed reachability specifications that extend reachability verdicts to vulnerabilities not yet covered by human-verified specs.

New Specification field in alert details surfaces whether a reachability result was produced by an AI-generated or human-verified specification.

Redesigned Organization Switcher for Faster Navigation (2026-06-29) BLOG ↗ STABLE

Socket's organization switcher now surfaces Settings and orgs upfront, with a searchable All Organizations view.

└──▷ HOW TO FIND IT

Quickly jump between organizations when managing multiple clients or teams.

📍 In the Socket console, click the organization switcher in the top navigation, then type an organization name in the All Organizations search box to filter and select it instantly.

Redesigned organization switcher surfaces Settings and organizations at the top level for faster access.

New searchable All Organizations view for users belonging to multiple organizations.

Alert Titles Added to CSV Exports (2026-06-23) BLOG ↗ STABLE

Socket CSV exports now include a human-readable Alert Title column and more consistent Upgrade Version values.

Adds a human-readable Alert Title column to Alerts CSV exports, enabling direct CVE identification without cross-referencing the dashboard.

Populates Upgrade Version values more consistently in vulnerability alert CSV exports.

Add Firewall events data export (2026-06-16) BLOG ↗ STABLE

Socket Firewall now supports exporting event logs to external data sources via the new Data Export integration.

└──▷ **HOW TO FIND IT**

Route Firewall event logs to an external SIEM or data pipeline for retention and analysis.

📍 In the console, go to Settings › Data Export › Create Integration, then select Telemetry events as the data type.

WAS THIS USEFUL? 👍 👎

LavaMoat

Feedback ⓘ allow-scripts-v5.1.0 NOTES ↗

tools for sandboxing your dependency graph

allow-scripts gains a `check` command to detect stale allow-lists in CI without making changes.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch allow-scripts-v5.1.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout allow-scripts-v5.1.0
```

└──▷ **TRY IT**

Catch unapproved or newly-added install scripts in CI before they can run — fails with a warning instead of silently passing.

```
$ npx @lavamoat/allow-scripts check
```

Adds a `check` command that warns when the allow-scripts configuration is out of date, designed for use in CI pipelines without modifying any files.

WAS THIS USEFUL? 👍 👎

Cisco skill-scanner

Feedback ⓘ 2.0.12 NOTES ↗ CODE ↗

The skill-scanner is a CLI security scanner for AI agent skills.

skill-scanner 2.0.12 adds ASCII smuggling detection, jailbreak YARA rule, and repo-name propagation for scans.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.0.12 https://github.com/cisco-ai-defense/skill-scanner.git
# already have the repo? check out this version:
$ git checkout 2.0.12
```

└──▷ TRY IT

Scan a skill for ASCII smuggling and jailbreak patterns — both new detections fire automatically via the static engine, no extra flags needed.

```
$ skill-scanner scan /path/to/skill --format sarif --output results.sarif
```

Scan a GitHub repo and surface jailbreak/ASCII-smuggling findings in CI, with the repo name now correctly propagated in scan metadata.

```
$ skill-scanner scan-repo owner/repo --use-llm --fail-on-severity high --format sarif --output results.sarif
```

Adds detection of ASCII smuggling via the Unicode Tag Block, catching a covert data-hiding technique used in prompt injection attacks.

Adds a jailbreak detection YARA rule to the static analysis engine.

Extends scanner to pass repository name through scan context, improving traceability in multi-repo pipelines.

Supports the OpenAI `user` request field in LLM analyzer calls.

Enables temperature override and automatic temperature omission for newer Claude and o1 models in the LLM analyzer.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian SaaS Feedback ⓘ

Generic Detectors landing on your non-VCS sources, powered by the False Positive Remover

3 RELEASES · 2026-06-19 → 2026-07-07 BLOG ↗ LATEST STABLE

GitGuardian SaaS is a secrets detection and security platform that scans code repositories to prevent sensitive data leaks and policy violations.

GitGuardian extends Generic Password and High Entropy Secret detection to Confluence, Jira, Slack, and Teams — with AI false-positive filtering built in.

└──▷ HOW TO FIND IT

Review or toggle AI-powered false-positive filtering for non-code sources after it is enabled by default on Business accounts.

> 📍 In the GitGuardian console, go to Settings › AI settings and confirm that AI features (including the False Positive Remover) are enabled for your non-code source integrations.

Adds Generic Password and Generic High Entropy Secret detection to non-code sources including Confluence, Jira, Slack, Microsoft Teams, and others — matching the coverage already available for VCS.

Extends the False Positive Remover (LLM-based model) to non-code source findings, automatically tagging 25–40% of generic findings as false positives before they surface in the dashboard.

Adds remediation guidance content and detector-specific links as assignable metadata fields in the Jira Notifier integration ticket description template.

▾ 2 more releases in this issue · 2026-06-19 → 2026-07-07

ServiceNow Integration: Event Group Selection, Public Monitoring & Honeytoken Support

BLOG ↗ STABLE

GitGuardian's ServiceNow integration gains Public Monitoring, Honeytoken support, and granular event subscription filters.

└──▷ TRY IT

Retrieve the activity log for a specific secret incident via the Public API to feed into an automation or SIEM workflow.

```
$ curl -s -H 'Authorization: Token <your_api_token>'
'https://api.gitguardian.com/v1/incidents/<incident_id>/activity' |
jq .
```

Enable Honeytoken and Public Monitoring event routing in a new ServiceNow integration so decoy and perimeter alerts land in the same ticketing workflow as internal incidents.

> 📍 In the GitGuardian dashboard, go to Integrations › ServiceNow › Configure, then under `Event Groups` enable `Public monitoring` and `Honeytoken events`. In the `Notify when` section, select the additional events (e.g. Incident ignored, Incident resolved) to be posted as comments on the ticket.

Adds Event Group Selection to the ServiceNow integration: choose to route Internal Monitoring, Public Monitoring, or Honeytoken incidents into ServiceNow tickets.

Supports Public Monitoring in ServiceNow, so secrets exposed on public sources flow into the same remediation workflow as internal incidents.

Supports Honeytoken events in ServiceNow, bringing decoy-token alerts into the ticketing and remediation workflow.

New `Notify when` section enables granular event subscription: the New incident detected event creates the ticket, while additional selected events are posted as comments, keeping full incident history in one place.

Adds severity and risk score information to incident notifications in Slack and Microsoft Teams.

`+3 more`

`Smart Notifiers - filter alerts, tickets, and emails by incident criteria (2026-06-19)`

`BLOG ↗` `STABLE`

Smart Notifiers lets GitGuardian alert only on incidents that match severity, ML risk, validity, or secret type — no middleware needed.

└──▷ HOW TO FIND IT

Route only high-severity, valid secrets to a Slack channel — eliminating noise from informational or unverified incidents.

> 📍 In the console, go to Alerting & Notifications › Slack notifier settings, then add a filter rule: set Severity to `High` or `Critical` and Validity to `Valid`. Save the rule to apply it to that channel.

Prevent false-positive check run failures caused by secrets merged in from the target branch during large PRs.

📍 In the console, go to Integrations › GitHub › Check Runs settings and enable `Skip merge commits` to restrict flagging to secrets introduced on the PR branch only.

Unlink a runaway or mistakenly created Jira ticket from an incident without deleting the external ticket.

📍 In the console, open the incident detail page, locate the linked Jira ticket section, and click `Dissociate ticket` to remove the association while keeping the Jira issue intact.

Adds Smart Notifiers: per-channel filtering rules for webhooks, Slack, and Microsoft Teams alerts based on severity, ML risk score, validity, secret type, or GitGuardian tag across internal and public monitoring.

Enables email alerting to filter public monitoring incidents by severity.

Adds sources health management for Confluence Data Center and Microsoft Teams: auto-pauses ingestion on unreachable sources and auto-resumes when health is restored.

Expands Overview analytics to include Non-Human Identity (NHI) and Public Monitoring data.

New GitHub check run setting to skip merge commits in pull request scans, flagging only secrets introduced on the PR branch.

`+6 more`

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine

Feedback ⓘ

Detection Engine Updates Version 2.166 (2026-06-23) BLOG ↗

GitGuardian's Detection Engine Release Notes documents updates to its automated secrets detection system that scans code repositories for exposed credentials and sensitive data.

GitGuardian adds 9 new secret detectors (Amadeus, Kinde, TomTom, Datadog, 5 Azure services) plus Gemfury validity checkers.

New detector for Amadeus OAuth Credentials.

New detector for Kinde OAuth Credentials.

New detector for TomTom API Key.

New detector for Datadog Personal Access Token.

New detector for Azure Maps Key.

+6 more

WAS THIS USEFUL? 👍 👎

GitGuardian Self-Hosted

Feedback ⓘ 2026.6 (2026-06-19)

BLOG ↗

GitGuardian Self-Hosted Release Notes is a changelog documenting updates, fixes, and new features for the self-hosted version of GitGuardian's secrets detection platform.

GitGuardian Self-Hosted 2026.6 adds 20+ new secret detectors, a premium scan retry worker, Teams granular alerting, and open-source AI agent skills.

└──▷ **USE IT**

Route failed large-repository VCS scans to a dedicated high-memory pod so base scanner pods stay small and only heavy retries scale up.

```yaml
# In your Helm values file, enable the premium scan retry worker:
premiumScanRetryWorker:
  enabled: true
  resources:
    requests:
      memory: "8Gi"
    limits:
      memory: "16Gi"
  replicaCount: 1
```

Receive only high-severity incident notifications in a dedicated Teams channel instead of all alert types.

📍 In the GitGuardian console, go to Integrations › Microsoft Teams, select a channel, then under `Alert types` enable only the specific incident event types (e.g. `Incident opened`, `Severity: critical`) relevant to that channel.

New premium scan retry worker (Helm only): routes failed large-repo VCS scans to a dedicated high-memory, scale-to-zero pod, keeping base scanner pods small.

New open-source agent-skills repository (GitGuardian/agent-skills) teaches AI coding assistants (Claude Code, Cursor) when to scan, how to read findings, and how to guide developers through remediation.

Secrets Detection Engine v2.163 adds 6 new detectors: Aikido API OAuth2 Credentials, Logflare Access Token, PolyAPI API Key, Apollo.io API Key, AES Cipher Key, and SAP OAuth Credentials.

Secrets Detection Engine v2.164 adds 6 new detectors: Authentik API Token, Pagar.me Encryption Key, Adobe Refresh Token, LiteLLM API Key, brapi.dev API Key, and Open

VSX Access Token, plus a new Google OAuth2 Keys checker and Adafruit IO API Key analyzer.

Secrets Detection Engine v2.165 adds 8 new detectors (Stitch API Key, ActiveMQ Credentials, Instantly API Key, Gitea Access Token, Fal.ai API Key, SurrealDB Cloud JWT Token, Kimi API Key, Wiz OAuth Credentials) and 8 new checkers including Okta, Azure, AWS Cognito, LaunchDarkly, Pulumi, and PyPI.

+4 more

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Cloudflare

Feedback ⓘ

WAF - WAF Release - Scheduled changes for 2026-07-13 (2026-07-06)

7 RELEASES · 2026-06-15 → 2026-07-06 BLOG ↗ LATEST STABLE

Cloudflare Security Changelog is an RSS feed that tracks and announces updates, changes, and new features related to Cloudflare's application security products.

WAF adds new detections for Citrix NetScaler ADC and Progress Kemp LoadMaster RCE vulnerabilities.

New detection for Citrix NetScaler ADC insufficient input validation (CVE-2026-8451), rule 78826e3223b94da493a2ade876973ac4, deploying in Log mode on 2026-07-13.

New detection for Progress Kemp LoadMaster remote code execution (CVE-2026-8037), rule 6b64d216620449fbb273d07910233f36, deploying in Log mode on 2026-07-13.

▾ 6 more releases in this issue · 2026-06-15 → 2026-07-06

Bots - More visibility into bot traffic with BotBase and Attribution Business Insights

BLOG ↗ STABLE

Enterprise Bot Management gains BotBase bot directory and Attribution Business Insights crawl-to-referral dashboard.

.

└──▷ HOW TO FIND IT

Use a bot's detection ID from BotBase to build a precise Security rule that blocks a specific crawler on your zone.

> 📍 In the Cloudflare dashboard, go to Security › Bots › BotBase, search for the target bot, copy its detection ID, then navigate to Security › WAF › Custom Rules and create a rule matching that detection ID.

Review which AI crawlers are sending referral traffic back versus just consuming your content, to inform allow/block decisions.

> 📍 In the Cloudflare dashboard, go to Security › Bots › Attribution Business Insights, set the time range to 30 days, and sort operators by crawl-to-referral ratio to identify high-crawl/low-referral operators.

Adds BotBase, a searchable in-dashboard directory of every bot Cloudflare tracks, classified by behavior (Search, Agent, Training, Transact, Data Collection, SEO, Ads

Verification).

Enables filtering live zone traffic down to a single bot within BotBase to investigate its activity and copy its detection ID for use in Security rules.

Introduces Attribution Business Insights dashboard showing crawl-to-referral ratios per bot operator over 24 hours, 7 days, or 30 days, with Allowed/Blocked/Partially blocked action status.

Publishes every tracked bot from BotBase to Cloudflare Radar's public bots and agents directory.

`Bots - New options to manage AI traffic (2026-07-01)` BLOG ↗ STABLE

Cloudflare now lets all plans block AI crawlers by behavior — Search, Agent, or Training — independently per page type.

.

└──▷ **HOW TO FIND IT**

Opt out of the September 15, 2026 default changes before they apply to your newly onboarded domain.

> 📍 In the Cloudflare dashboard, go to Security › Bots › AI Traffic and manually configure each behavior preset (Search, Agent, Training) to your preferred setting before September 15, 2026.

Introduces three independently controllable AI traffic behavior categories — Search, Agent, and Training — available to all customers including Free plan.

Enables per-behavior blocking policies: block on all pages, block only on ad-displaying pages, or allow — giving fine-grained control over which AI crawlers can access content.

New default policy (effective September 15, 2026) automatically blocks Training and Agent crawlers on ad-displaying pages for newly onboarded domains, with opt-out available before that date.

Multi-purpose crawlers combining Search and Training behaviors will be subject to Training-class blocking rules under the new defaults.

`WAF - WAF Release - 2026-07-01 (2026-07-01)` BLOG ↗ STABLE

WAF adds FortiSandbox path traversal detection (CVE-2026-39813) and disables Fake Bing/MSN Bot rule.

└──▷ **HOW TO FIND IT**

Override the new FortiSandbox rule from Log to Block immediately if you run FortiSandbox and want active protection rather than waiting for Cloudflare's default escalation.

📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 32075e19b1494117ac5915e8d84c92c9, and set the action override to Block.

Re-enable blocking of fake Bing/MSN bots via a custom WAF rule if your zone depended on rule ae20608d93b94e97988db1bbc12cf9c8 for bot enforcement.

📍 In the Cloudflare dashboard, go to Security › WAF › Custom Rules, create a new rule with field User Agent matching known fake-bot patterns, and set action to Block.

New detection for Fortinet FortiSandbox path traversal (CVE-2026-39813): rule 32075e19b1494117ac5915e8d84c92c9 blocks unauthenticated arbitrary file reads, currently set to Log action.

Breaking change: Anomaly:Header:User-Agent - Fake Bing or MSN Bot rule (ae20608d93b94e97988db1bbc12cf9c8) action changed from Enabled/Block to Disabled — zones relying on this rule for bot enforcement must now implement alternative coverage.

`WAF - WAF Release - 2026-06-23 (2026-06-23)` BLOG ↗ STABLE

New managed rule blocks CVE-2026-10520 pre-auth OS command injection in Ivanti Sentry.

└──▷ **HOW TO FIND IT**

Verify the new Ivanti Sentry rule is active in Block mode on your zone — useful immediately after deployment to confirm coverage before patching.

📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 500a90789f874345b60b0de7242fdf83, and confirm the action is set to Block.

`WAF - Use Cloudforce One threat intelligence in WAF rules (2026-06-15)` BLOG ↗

STABLE

WAF custom rules can now match Cloudforce One threat intelligence fields to block IPs flagged for recent attack activity.

└──▷ **TRY IT**

Block inbound requests from IPs linked to DDoS campaigns that have targeted France — useful for protecting French-market infrastructure during active threat periods.

```
$ any(cf.intel.ip.target_countries[*] == "FR") and
any(cf.intel.ip.datasets[*] == "ddos")
```

Adds `cf.intel.ip.*` fields populated from Cloudforce One threat intelligence, enabling WAF custom and rate-limiting rules to act on IPs flagged within the past seven days.

New `cf.intel.ip.datasets` field identifies whether a flagged IP was associated with DDoS or WAF attack activity.

New `cf.intel.ip.attacker_names` and `cf.intel.ip.attacker_countries` fields expose known threat actor identities and source countries for use in rule expressions.

New `cf.intel.ip.target_industries` and `cf.intel.ip.target_countries` fields allow rules to prioritize defense based on the sectors or regions an IP has historically targeted.

Threat intelligence matches are recorded in Security Analytics and fully supported via the Cloudflare API and Terraform.

WAF - WAF Release - 2026-06-15 (2026-06-15) BLOG ↗ STABLE

WAF adds Ghost CMS blind SQLi detection (CVE-2026-26980) and a new generic obfuscated boolean SQLi bypass rule.

└──▷ **HOW TO FIND IT**

Enable the new obfuscated boolean SQLi rule (shipped as Disabled) to actively block evasion attempts against your zone.

> 📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules › Cloudflare Managed Ruleset, search for rule ID 6c64b68ef5ed45e7a622cdaab56f403f (SQLi - Obfuscated Boolean - URI), and change its action from Disabled to Block.

New managed rule 439c4ef64b32447989bdf412b4c29bc6 detects and blocks blind SQL injection in Ghost CMS Content API (CVE-2026-26980, versions 3.24.0–6.19.0), moving to Block action.

New managed rule 6c64b68ef5ed45e7a622cdaab56f403f identifies generic SQLi bypass attempts using obfuscated boolean logic in URI parameters (currently set to Disabled).

WAS THIS USEFUL? 👍 👎

Arcjet Feedback ⓘ

SDK releases, security briefing improvements, new activity view (2026-07-03)

BLOG ↗

Arcjet is a developer-focused security SDK that protects web applications against bots, rate limiting, email validation, and attack threats with minimal code integration.

Arcjet adds proxy support, Nosecone Nonce API, enhanced security briefings, and a new Console activity view.

.

└──▷ **USE IT**

Inject a CSP nonce into a custom script tag from a Next.js Server Component when the <Script> component is unavailable.

javascript

```
import { nonce } from "@nosecone/next";

export default async function Page() {
  const cspNonce = await nonce();
  return <script nonce={cspNonce}>{`/* your inline script */`}
</script>;
}
```

Adds proxy support to the JS SDK 1.6.

Introduces an optimized Wasm bundle across JS, Go, and Python SDKs.

Delivers the first numbered Go SDK release (0.1.0) with improved per-request caching.

Adds a `nonce` function to the Nosecone Next.js library (v1.6.1), enabling nonce access from Server Components and Route Handlers.

Expands weekly security briefing emails with detailed threat tagging, country breakdowns, and updated recommendations.

+1 more

WAS THIS USEFUL? 👍 👎

Wallarm

Feedback ⓘ Introducing the New Attacks Interface (2026-07-02)

BLOG ↗

Wallarm is a cloud-native application security platform that protects APIs, microservices, and web apps from threats like injections, bots, and account takeovers.

Wallarm launches a redesigned Attacks interface with interactive dashboards, customizable views, and drill-down request inspection.

└──▷ HOW TO FIND IT

Scope an investigation to a specific application and attack type, then export the filtered results for an incident report.

> 📍 In the Wallarm console, go to Attacks › apply filters for the target application and attack type › set your time range › select columns relevant to the incident › click Export to download the tailored dataset.

Create a persistent dashboard for a specific host so the SOC team always has a scoped view ready without re-applying filters each session.

> 📍 In the Wallarm console, go to Attacks › set filters for the target host and preferred columns and time range › click Save View › name the view and assign it as a dedicated dashboard for that host.

New interactive Attacks dashboard lets analysts navigate attack trends, spot anomalies, and apply filters directly from charts, tables, and statistics.

Customizable Views enable saving preferred columns, filters, time ranges, and groupings as dedicated dashboards scoped to specific attack types, applications, hosts, or incidents.

Attack Investigation Flow surfaces request-level details — source IP, host, URI, payload, parameters, classifications, headers, and full raw request — from a single drill-down.

Interactive Request Grouping lets analysts cluster attacking requests by attack type, source IP, host, path, or any custom combination to surface hard-to-spot patterns.

Customizable Export tailors attack data exports by filter, column selection, time range, and grouping for incident response, compliance, or offline analysis.

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ DataDome Terraform provider v2.4.0 (2026-06-30)

2 RELEASES · 2026-06-30 BLOG ↗ LATEST STABLE

DataDome is a bot protection and fraud prevention platform that detects and blocks malicious traffic, scraping, credential stuffing, and other automated threats in real time.

DataDome Terraform provider v2.4.0 adds rate limiting, time-box policies, agentic endpoint types, and new response actions for custom rules.

└──▷ **USE IT**

Protect an agentic login endpoint with a rate-limit policy to throttle automated credential-stuffing agents.

hcl

```hcl
resource "datadome_custom_rule" "agentic_login_rate_limit" {
  name          = "Rate-limit agentic login"
  endpoint_type = "agentic-login"
  response      = "intent_based"
  policy        = "rate_limit"
}
```

Apply a time-box policy with a device_check response to block suspicious agentic cart activity within a rolling window.

hcl

```hcl
resource "datadome_custom_rule" "agentic_cart_timebox" {
  name          = "Time-box agentic cart abuse"
  endpoint_type = "agentic-cart"
  response      = "device_check"
  policy        = "time_box"
}
```

Adds `rate_limit` and `time_box` policy options for custom rules, enabling threshold-based and time-windowed enforcement.

Adds six new `endpoint_type` values for custom rules targeting agentic AI traffic: `agentic-general`, `agentic-account-creation`, `agentic-login`, `agentic-cart`, `agentic-forms`, and `agentic-payment`.

Adds three new `response` values for custom rules: `device_check`, `intent_based`, and `monetize`, expanding the range of enforcement actions available.

▾ **1 more release in this issue** · 2026-06-30

> iOS SDK for Apollo 4.0.0 (2026-06-30) BLOG ↗ STABLE
>
> DataDome iOS SDK for Apollo 4.0.0 adds support for Apollo v2.
>
> Supports Apollo v2 GraphQL client for iOS, enabling DataDome bot protection on Apollo v2-based apps.

WAS THIS USEFUL?  

Invicti

Feedback ⓘ Release 20260625 (2026-06-25) BLOG ↗

Invicti is an automated web application security scanner that identifies vulnerabilities like SQL injection and XSS to help organizations secure their web assets.

Invicti adds gRPC scanning, bulk scan suspension, collection dashboards, and richer CSV exports in this release.

└─▷ **HOW TO FIND IT**

Add gRPC DAST coverage to a microservice by supplying its Protobuf contract at scan configuration time.

> 📍 In the Invicti console, go to New Scan › API Scanning, select gRPC as the API type, and upload the service's .proto definition file before launching the scan.

Pause all automated scanning org-wide during a change-freeze window without touching each schedule individually.

> 📍 In the Invicti console, go to Scheduled Scans and use the Suspend All Scheduled Scans control to halt every future scheduled scan across the account at once.

Reduce scan failures on sites with CAPTCHA or MFA when scanning through an internal agent by enabling interactive login.

> 📍 In the scan profile settings for an internal agent scan, enable the AI-aided Login setting so Invicti can use AI to recover failed Login Sequence Recorder playbacks automatically.

Supports scanning of gRPC services by uploading .proto definition files, extending DAST coverage to gRPC-based APIs.

New Suspend all scheduled scans control lets organization administrators pause all future scheduled scans across the account in one action.

Adds a dedicated dashboard for Collections, providing at-a-glance security posture visibility at the collection level.

New AI-aided Login setting enables AI recovery when Login Sequence Recorder (LSR) playback fails during a scan, reducing authentication-related scan failures.

Extends auditing to SCIM configuration changes, recording who changed provisioning settings and when in the audit log.

+3 more

WAS THIS USEFUL? 👍 👎

Acunetix 360 Feedback ⓘ v26.6.1 – 16 June 2026 (2026-06-16)

BLOG ↗

Acunetix 360 On-Demand Changelog is an RSS feed tracking updates, fixes, and new features released for the Acunetix 360 cloud-based web application security scanner.

Acunetix 360 v26.6.1 adds pre-scan auth validation, interactive MFA login, and sitemap data retention.

New pre-scan authentication validation checks that credentials are working before a full scan begins.

New interactive MFA login support enables authenticated scanning of targets protected by multi-factor authentication.

New sitemap data retention preserves discovered sitemap information across scans for ongoing analysis.

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite Feedback ⓘ

DAST 2026.6 (2026-06-16) BLOG ↗

Burp Suite is a leading web application security testing platform used by professionals to find, exploit, and report vulnerabilities in web applications.

Burp Suite DAST 2026.6 adds wildcard and IPv6 support to proxy exclusion lists and a new scanner update filename format.

Supports wildcard entries in the proxy exclusion list for flexible traffic filtering.

Supports IPv6 addresses in the proxy exclusion list.

Supports the new scanner update filename format.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

iLEAPP

Feedback ⓘ v2026.1.0 NOTES ↗ CODE ↗

iOS Logs, Events, And Plist Parser

iLEAPP v2026.1.0 adds 10+ new parsers, conversation views, custom artifact paths, and broad LAVA migration.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.1.0 https://github.com/abrignoni/iLEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.1.0
```

└──▷ TRY IT

Load custom or case-specific artifact modules from a separate folder without touching the iLEAPP install — useful for proprietary parsers or in-progress modules.

```
$ ileapp -t zip -i /cases/iphone_extraction.zip -o /cases/output/ --
custom_artifacts_path /cases/my_custom_artifacts/
```

Name the report subfolder after a case or device for cleaner multi-device output directories.

```
$ ileapp -t fs -i /mnt/extraction/ -o /cases/output/ --
custom_output_folder iPhone13_JohnDoe_2026
```

Parse an encrypted iTunes backup by supplying the backup password directly on the CLI for scripted/automated workflows.

```
$ ileapp -t itunes -i /cases/backup/ -o /cases/output/ --
itunes_password S3cr3tPassw0rd
```

Adds new parsers for Zalo Chat and Users, Discord cache, Home Depot app, Oura Ring, App Web Usage (Biome), Biome NotificationUsage, springboard wallpaper, Apple ATX image files, iOS Widgets, and keychain artifacts.

Adds conversation views to WhatsApp, Viber, Kik, Slack, Life360, TikTok, Discord, Teleguard, Reddit, LinkedIn, and Telegram.

Adds `--custom_artifacts_path` CLI flag to load artifact modules from an extra folder, enabling third-party or case-specific parsers without modifying the install.

Adds `--custom_output_folder` CLI flag to set a custom name for the report output subfolder.

Adds an App Inventory module for parsing-coverage analysis across installed apps.

+9 more

WAS THIS USEFUL?  

Obsidian Forensics Hindsight

Feedback ⓘ v2026.06

NOTES ↗ CODE ↗

Browser forensics tool for Google Chrome (and other Chromium-based browsers)

Hindsight v2026.06 adds full Mozilla Firefox support and expands Chrome extension, download, and Service Worker parsing.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.06
https://github.com/RyanDFIR/hindsight.git
# already have the repo? check out this version:
$ git checkout v2026.06
```

.

└──▷ TRY IT

Analyze a Firefox profile on Windows and export results to XLSX for timeline review.

```
$ hindsight.py -i
"C:\Users\Investigator\AppData\Roaming\Mozilla\Firefox\Profiles\abc123
.default" -b Firefox -o firefox_case
```

Analyze a Chrome profile with all expanded download metadata and extension details written to XLSX.

```
$ hindsight.py -i "C:\Users\Suspect\AppData\Local\Google\Chrome\User
Data\Default" -o chrome_case -f XLSX
```

Adds Mozilla Firefox browser support: parses places.sqlite (history, bookmarks, downloads), cookies.sqlite, formhistory.sqlite, permissions.sqlite, logins.json, extensions.json, sessionstore, favicons, cache2, localStorage, IndexedDB, and bounce-tracking/content-blocking databases.

Adds Firefox HSTS pin parsing from `SiteSecurityServiceState.bin` and forensically relevant preferences from prefs.js.

Adds Firefox bookmark backup parsing from bookmarkbackups/*.jsonlz4, enabling recovery of deleted bookmarks from daily snapshots.

Adds parsing of Service Worker-related data, including Background Fetch, Background Sync, Cache Storage, Content Index, Cookie Store, Notifications, and Payments

protobuf artifacts.

Enhances Chrome extension handling: merges data from disk and Secure Preferences, adds a dedicated Extensions worksheet in Excel output, and surfaces detailed permissions and content scripts.

+4 more

WAS THIS USEFUL? 👍 👎

Yamato Security Hayabusa

Feedback ⓘ v3.10.0 NOTES ↗

CODE ↗

Hayabusa (隼) is a sigma-based threat hunting and fast forensics timeline generator for Windows event logs.

Hayabusa v3.10.0 adds Sigma 2.1 `|neq` field modifiers and per-tactic unique/total alert counts in MITRE ATT&CK output.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.10.0 https://github.com/Yamato-Security/hayabusa.git
# already have the repo? check out this version:
$ git checkout v3.10.0
```

Supports Sigma 2.1 `|neq` and `|fieldref|neq` field modifiers, enabling negated comparisons composable with `|contains`, `|startswith`, `|endswith`, and `|fieldref` in detection rules.

Adds unique and total alert counts per MITRE ATT&CK tactic in detection results, giving practitioners a clearer picture of tactic coverage across a scan.

WAS THIS USEFUL? 👍 👎

Request Tracker / RTIR (Best Practical)

Feedback ⓘ

No RT Expertise Required: Meet the Queue Creation Assistant (2026-07-02) BLOG ↗

Request Tracker (RT) and RTIR are open-source ticketing and incident response systems that help organizations track, manage, and resolve requests and security incidents.

RT-Extension-AI adds a Queue Creation Assistant that builds queues, workflows, custom fields, groups, and rights from plain-language conversation.

New Queue Creation Assistant in RT-Extension-AI lets admins describe a queue's purpose in plain language and generates the full configuration — lifecycle/workflow, custom fields, groups, rights, and watchers — for review before applying.

Conversational setup flow replaces manual multi-step queue configuration: the assistant asks targeted questions about team, ticket types, and workflow, then presents a complete plan for approval before making any changes.

One-step queue provisioning: after admin approval, the assistant creates the queue and all supporting elements (workflow, custom fields, groups, rights, watchers) atomically in a single action.

WAS THIS USEFUL? 👍 👎

Google Timesketch

Feedback ⓘ 20260630

2 RELEASES · 2026-06-17 → 2026-06-30 NOTES ↗ LATEST STABLE

Collaborative forensic timeline analysis

Timesketch 20260630 adds SecGemini analyzer results to report summaries and enriches OpenTelemetry spans with search metadata.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 20260630 https://github.com/google/timesketch.git
# already have the repo? check out this version:
$ git checkout 20260630
```

Adds SecGemini analyzer report output to the sketch report summary, surfacing AI-driven analysis alongside other analyzer results.

Updates SecGemini provider to align with the sec-gemini v3 SDK APIs, enabling compatibility with the latest SecGemini capabilities.

Adds search metadata attributes to OpenTelemetry spans, enabling deeper observability into search query telemetry.

▾ 1 more release in this issue · 2026-06-17 → 2026-06-30

20260617 NOTES ↗ STABLE

Timesketch 20260617 enforces cross-sketch ownership on event fact conclusions via the API.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 20260617
https://github.com/google/timesketch.git
# already have the repo? check out this version:
$ git checkout 20260617
```

└──▷ BREAKING ON UPGRADE

! OpenSearch users must upgrade to at least version 3.7.0 before upgrading to this Timesketch release if running OpenSearch >= 3.x

WAS THIS USEFUL?  

Microsoft AVML

Feedback ⓘ v0.19.0 NOTES ↗ CODE ↗

AVML - Acquire Volatile Memory for Linux

AVML v0.19.0 collapses all subcommands into one binary and adds a `stream` subcommand for diskless memory capture.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.19.0 https://github.com/microsoft/avml.git
# already have the repo? check out this version:
$ git checkout v0.19.0
```

└──▷ **TRY IT**

Capture memory from a live host directly to Azure Blob Storage with no local file written — useful when disk space is scarce or forensic integrity requires avoiding local writes.

```
$ avml stream blob ${SAS_URL}
```

Stream a memory snapshot over TCP to a waiting collector — pipe straight into a file on the analyst's machine without touching the target's disk.

```
$ nc -l 9000 > snapshot.lime
# then on the target:
avml stream tcp collector.example.com:9000
```

New `stream` subcommand streams a live memory snapshot directly to Azure Block Blob or a remote TCP listener without writing to local disk — ideal for read-only roots, limited-disk hosts, or forensic chain-of-custody requirements.

New `stream blob <SAS_URL>` destination pipes memory directly to Azure Block Blob Storage, auto-sizing block size to stay within Azure's 50,000-block limit; `--sas-block-size` (MiB) sets a floor and `--sas-block-concurrency` caps in-flight uploads.

New `stream tcp <host:port>` destination streams a snapshot over a raw TCP connection to a netcat-style collector.

All subcommands (`acquire`, `convert`, `upload`, `stream`) are now consolidated into a single `avml` binary instead of separate binaries.

Enforces that `--url` and `--sas-url` are mutually exclusive on `avml acquire`, and requires one of them when `--delete` is set — preventing silent misconfiguration.

└──▷ **BREAKING ON UPGRADE**

! The `avml-upload` binary has been collapsed into `avml`; the separate `avml-upload` binary no longer exists.

! `--url` and `--sas-url` are now mutually exclusive on `avml acquire`; previously both could be specified together.

! `--delete` now requires `--url` or `--sas-url` to be set; previously `--delete` could be used without a remote URL.

WAS THIS USEFUL? 👍 👎

ElcomSoft Feedback ⓘ

iOS Forensic Toolkit 10.10 adds pairing-free sideloading of the extraction agent (2026-0

2 RELEASES · 2026-06-18 → 2026-06-24 BLOG ↗ LATEST STABLE

ElcomSoft Product News is an RSS feed delivering updates about ElcomSoft's password recovery, forensic, and security auditing software products and releases.

iOS Forensic Toolkit 10.10 enables pairing-free sideloading of the extraction agent via wired or wireless network bridge.

Supports pairing-free sideloading of the low-level extraction agent, eliminating the need to pair the device before delivery.

Delivers the extraction agent over a wired or wireless network bridge as alternative sideloading methods.

▾ **1 more release in this issue** · 2026-06-18 → 2026-06-24

> Elcomsoft Phone Breaker 11.2: downloads iOS 26 iCloud backups (2026-06-18)
>
> BLOG ↗ STABLE
>
> Elcomsoft Phone Breaker 11.2 adds support for downloading iCloud backups from iOS/iPadOS 26 and iOS 27 beta.
>
> Supports downloading iCloud backups created on iOS 26 and iPadOS 26 devices.
>
> Supports downloading iCloud backups from iOS/iPadOS 27 beta builds.

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

Datadog GuardDog

Feedback ⓘ v3.1.0

3 RELEASES · 2026-06-16 → 2026-07-08 NOTES ↗ CODE ↗ LATEST STABLE

GuardDog is a CLI tool to Identify malicious PyPI and npm packages

GuardDog v3.1.0 adds a PyPI Inspector link for remote findings and a new npm risky_new_dependency metadata rule.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.1.0 https://github.com/DataDog/guarddog.git
# already have the repo? check out this version:
$ git checkout v3.1.0
```

Displays a PyPI Inspector URL in human-readable output for remote PyPI scans that have findings, making file-level triage faster.

Adds new npm metadata rule `risky_new_dependency` that flags packages introducing a risky new transitive dependency.

▾ 2 more releases in this issue · 2026-06-16 → 2026-07-08

v3.0.0 NOTES ↗ CODE ↗ STABLE

GuardDog v3 adds kernel-level sandboxing, S3 scanning, risk-based scoring engine, and removes legacy CLI commands.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.0.0 https://github.com/DataDog/guarddog.git
# already have the repo? check out this version:
$ git checkout v3.0.0
```

└──▷ TRY IT

Scan an untrusted PyPI package with the kernel-level sandbox enforced (default), or explicitly disable it on a platform where Landlock/Seatbelt is unavailable.

```
$ guarddog pypi scan requests
guarddog pypi scan /tmp/suspicious.tar.gz --no-sandbox
```

Scan a private npm package stored in an S3 bucket — useful for auditing internally mirrored or pre-release packages before deployment.

```
$ guarddog npm scan s3://my-bucket/path/to/package/
guarddog pypi scan s3://my-bucket/path/to/package.tar.gz
```

Scan a local package archive and supply a metadata JSON sidecar so metadata-based rules (e.g. typosquatting, maintainer checks) also fire.

```
$ guarddog pypi scan /tmp/triage.tar.gz --metadata /tmp/triage-
metadata.json --output-format json
```

Adds kernel-level sandbox (Linux via Landlock, macOS via Seatbelt) for all scan operations — blocks network access and restricts filesystem during archive extraction and YARA analysis; use `--no-sandbox` to opt out explicitly.

Supports scanning packages stored in S3 via `s3://` URIs, with AWS credential verification before scan and path-traversal-safe download to a temp directory.

Introduces a risk-based detection engine that correlates code capabilities with threat indicators in the same file to produce a 0–10 severity score, reducing false-positive alert fatigue.

Adds `--metadata` flag to supply a package metadata JSON file for local scans, enabling metadata rules that would otherwise require a registry lookup.

Adds `--zip-password` flag to scan encrypted ZIP archives (`.zip` , `.whl` , `.egg`), with `-` to read the password from stdin.

+1 more

└─▷ **BREAKING ON UPGRADE**

! The deprecated CLI commands `scan` and `verify` have been removed.

v3.0.0a2 NOTES ↗ CODE ↗ PRE-RELEASE

GuardDog v3.0.0a2 adds encrypted ZIP scanning, S3/remote archive support, and removes legacy top-level commands.

└─▷ **GET THIS VERSION** ▾

```
$ git clone --branch v3.0.0a2
https://github.com/DataDog/guarddog.git
# already have the repo? check out this version:
$ git checkout v3.0.0a2
```

└──▷ **TRY IT**

Scan a password-protected ZIP archive from a remote URL — useful for triaging malware samples stored in password-locked archives.

```
$ guarddog npm scan https://github.com/DataDog/malicious-software-
packages-
dataset/raw/refs/heads/main/samples/pypi/compromised_lib/mistralai/2
.4.6/2026-05-12-mistralai-v2.4.6.zip --zip-password infected
```

Scan a local encrypted wheel file, reading the password from stdin to avoid exposing it in shell history.

```
$ echo 'mypassword' | guarddog pypi scan /tmp/suspicious_package.whl
--zip-password -
```

Adds `--zip-password` flag to `scan` commands for extracting password-protected ZIP, WHL, and EGG archives (pass `-` to read from stdin).

Supports scanning remote archives via URL (e.g., GitHub raw links) in addition to local paths and S3 prefixes.

Breaking change: removes the deprecated top-level `scan` and `verify` commands — use ecosystem-prefixed commands (`pypi scan` , `npm scan` , etc.) instead.

└──▷ **BREAKING ON UPGRADE**

! The deprecated CLI commands `scan` and `verify` have been removed.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

CHIPSEC

Feedback ⓘ 2.0.6 2 RELEASES · 2026-06-16 → 2026-07-01

NOTES ↗ CODE ↗ LATEST STABLE

Platform Security Assessment Framework

CHIPSEC 2.0.6 restores the `uefi types` command and adds platform autodetection to `uefi tables`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.0.6 https://github.com/chipsec/chipsec.git
# already have the repo? check out this version:
$ git checkout 2.0.6
```

└──▷ TRY IT

Enumerate UEFI variable types on the running system — platform is now detected automatically, no manual specification needed.

```
$ chipsec_util uefi tables
```

List supported UEFI types on the current platform without manually identifying it first.

```
$ chipsec_util uefi types
```

Adds platform autodetection to `chipsec_util uefi tables`, removing the need to specify the platform manually.

Adds a decode dictionary to `ExitCode` in `cfg_checker.py`, making exit code interpretation programmatic.

Expands platform DID matching to HostCtrl DIDs for ADL, ARL, MTL, RPL, and TGLU platforms.

▾ 1 more release in this issue · 2026-06-16 → 2026-07-01

> 2.0.5 NOTES ↗ CODE ↗ STABLE
>
> CHIPSEC 2.0.5 adds Intel 440FX (QEMU) config, Meteor Lake DIDs, and expanded cfg_checker validation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.0.5 https://github.com/chipsec/chipsec.git
# already have the repo? check out this version:
$ git checkout 2.0.5
```

Adds platform configuration for Intel 440FX (QEMU), enabling CHIPSEC analysis against QEMU-emulated environments.

Adds Meteor Lake (MTL) Device IDs from the datasheet, extending hardware coverage for 14th Gen Intel Core platforms.

Adds BAR register/field validity and subcomponent type checks to cfg_checker, improving configuration correctness validation.

Splits tcobar and smbus_base into their own dedicated configuration files for cleaner platform config organization.

WAS THIS USEFUL? 👍 👎

checksec.sh

Feedback ⓘ 3.2.0 NOTES ↗ CODE ↗

Checksec

checksec 3.2.0 adds SafeStack, sanitizer, W^X, Selfrando, and stack-clash detection plus a `--fail-if` CI gate and CSV output.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 3.2.0 https://github.com/slimm609/checksec.git
# already have the repo? check out this version:
$ git checkout 3.2.0
```

└──▷ **TRY IT**

Gate a CI pipeline on RELRO, canary, and PIE — fail the build if any binary in a directory is missing these hardening properties.

```
$ checksec dir ./build/bin --recursive --fail-if relro,canary,pie
```

Audit a single binary for the full new check suite (SafeStack, sanitizers, W^X, stack-clash) and emit machine-readable CSV for ingestion into a spreadsheet or SIEM.

```
$ checksec file /usr/bin/myapp --output csv
```

Run a clean, warning-free scan of a directory with forced color off — useful in log-heavy CI environments where ANSI codes break parsers.

```
$ checksec dir /opt/app --recursive --no-warnings --color never --
output json
```

Adds SafeStack detection for ELF binaries by scanning for the `__safestack_init` symbol.

Extends CFI hardening detection to i386 CET and RISC-V Zicfilp/Zicfiss (landing pads and shadow stack), covering both RV32 and RV64.

New sanitizer detection (ASAN, UBSAN, MSAN, TSAN, LSAN) via compiler-rt symbol prefix scan.

New Selfrando detection via the `.txtrp` ELF section.

New W^X segment check (`SeparateCode`): flags any `PT_LOAD` segment that is both writable and executable.

+6 more

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ CLI v15.4.0 (2026-07-08)

16 RELEASES · 2026-06-16 → 2026-07-08 BLOG ↗ LATEST STABLE

Fastly is a cloud edge platform that accelerates websites and APIs, provides CDN services, and offers security features like DDoS protection and WAF.

Fastly CLI v15.4.0 adds compute install-tools command and NGWAF account/workspace timeseries support.

└──▷ TRY IT

Pre-install the Viceroy local testing binary before starting Compute development, useful in CI or fresh dev environments.

```
$ fastly compute install-tools
```

New `compute install-tools` command pre-installs the Viceroy binary for local Compute development environments.

New NGWAF `timeseries` subcommands support account-level and workspace-level time series data queries.

▾ 15 more releases in this issue · 2026-06-16 → 2026-07-08

Next-Gen WAF agent 4.80.1 (2026-07-08) BLOG ↗ STABLE

Next-Gen WAF agent 4.80.1 improves CMDEXE, path traversal, and XSS detection with refreshed GeoIP data.

Improves CMDEXE detection for more accurate command-injection identification.

Improves TRAVERSAL detection to better catch path-traversal attack patterns.

Improves XSS detection for broader cross-site scripting coverage.

Updates base GeoIP data to July 2026 for current geolocation accuracy.

Add Compute backend error metrics to Historical Stats and Real-Time Analytics APIs (20

BLOG ↗ STABLE

Fastly adds 16 backend Compute error metrics to Historical Stats and Real-Time Analytics APIs for root-cause troubleshooting.

└──▷ TRY IT

Query the Real-Time Analytics API to surface backend Compute error counts broken down by category (TLS, HTTP, DNS, connection) for a service in the last 60 seconds.

```
$ curl -s -H 'Fastly-Key: <FASTLY_API_TOKEN>'
'https://rt.fastly.com/v1/channel/<SERVICE_ID>/ts/h'
```

Add Notification Service Event Mappings API (2026-06-30) BLOG ↗ STABLE

Fastly adds Event Mappings API to control which audit events trigger notifications per service or workspace.

└──▷ **TRY IT**

Discover which audit event types are available and what scope types each supports before building a mapping.

```
$ curl -X GET https://api.fastly.com/notifications/event-types \
  -H 'Fastly-Key: <FASTLY_API_TOKEN>' \
  -H 'Accept: application/json'
```

Discover available scope categories (account, VCL service, Wasm service, WAF workspace) to use when scoping a new event mapping.

```
$ curl -X GET https://api.fastly.com/notifications/scope-types \
  -H 'Fastly-Key: <FASTLY_API_TOKEN>' \
  -H 'Accept: application/json'
```

New Event Mappings API enables creating, managing, and querying which Fastly audit events trigger notifications to configured integrations.

Supports scoping notification mappings account-wide or to specific CDN (VCL) services, Compute (Wasm) services, or Next-Gen WAF workspaces.

New discovery endpoint `/event-types` returns all supported audit event types with their compatible scope types.

New discovery endpoint `/scope-types` returns the supported scope categories for notification mappings.

Add new Routing Configurations API and associated support for it in the Domains API (2

BLOG ↗ STABLE

Fastly adds Routing Configurations API to route traffic by URL path/condition, with Domain API support for multi-domain rule binding.

New Routing Configurations API enables defining routing rules based on URL paths and conditions to direct traffic to different services.

Updated Domains API supports associating domains with routing configurations, enabling consistent rule application across multiple domains.

Connect your Fastly account to your Next-Gen WAF corp (2026-06-30) BLOG ↗

STABLE

Fastly accounts can now be linked to an existing Next-Gen WAF corp for unified control panel and API access.

└──▷ **HOW TO FIND IT**

When your team manages WAF rules in the Next-Gen WAF control panel but wants to administer them through the Fastly control panel and API going forward.

📍 In the Fastly control panel, follow the 'Connecting a Fastly account to your corp' guide to link your Fastly account to your existing Next-Gen WAF corp.

Enables connecting a Fastly account to an existing Next-Gen WAF corp, surfacing Next-Gen WAF and Bot Management through the Fastly control panel and Fastly API without losing access to the standalone Next-Gen WAF control panel or API.

WAF configurations and user accounts are shared and synchronized in real time across both control panels and both APIs after the account is connected.

New integration types for Notifications API (2026-06-30) BLOG ↗ STABLE

Fastly Notifications API adds Datadog, Jira, OpsGenie, and Splunk On-Call as integration targets.

Adds five new Notifications API integration types: Datadog, Jira Issue, Jira Service Management, OpsGenie, and Splunk On-Call.

Enables routing of observability alerts and audit log event notifications directly to incident management and issue tracking platforms.

CLI v15.3.0 (2026-06-26) BLOG ↗ STABLE

Fastly CLI v15.3.0 adds prefix-based bulk KV store deletion and contentguard bot management support.

└──▷ **TRY IT**

Bulk-delete all KV store keys sharing a common prefix — useful for purging session data, cache entries, or tenant records in one operation.

```
$ fastly kvstore-entry delete --store-id <store-id> --key-prefix
<prefix>
```

Adds support for deleting multiple KV store entries at once using a key prefix.

Supports `contentguard` configuration on the `bot_management` product under [setup.products] in compute deploy.

`Account-level requests for Fastly control panel (2026-06-25)` **BLOG ↗** **STABLE**

Fastly control panel now shows account-level request visibility across all workspaces on the Requests page.

└──▷ **HOW TO FIND IT**

View and filter requests across every workspace in your account to identify traffic patterns or anomalies without switching between individual workspaces.

📍 In the Fastly control panel, go to Requests and select the account-level view to filter requests across all workspaces.

`Go SDK 1.8.0 (2026-06-23)` **BLOG ↗** **STABLE**

Fastly Go SDK 1.8.0 adds bot detection, VPN/proxy detection, stale-if-error caching, and richer KV store controls.

Adds bot detection hostcalls via `fsthttp`, enabling Compute apps to identify bot traffic at the edge.

Adds `ResVPNProxy` support in `fsthttp` to detect and act on requests originating from VPNs or proxies.

Adds stale-if-error logic to the HTTP caching layer, allowing cached responses to be served when origin errors occur.

Adds `IfGenerationMatch` to `InsertConfig` in `kvstore`, enabling conditional writes based on KV store entry generation.

Supports streaming request bodies on KV store insert operations.

`+1 more`

`Next-Gen WAF: Add account-level requests endpoint (2026-06-22)` **BLOG ↗**

STABLE

Next-Gen WAF adds account-level requests endpoint to search across all workspaces in one call.

└──▷ **TRY IT**

Query WAF requests account-wide to hunt for a suspicious IP pattern without scoping to a single workspace.

```
$ curl -X GET 'https://api.fastly.com/ngwaf/v1/requests?
q=ip:203.0.113.42' \
  -H 'Fastly-Key: <your-api-token>' \
  -H 'Accept: application/json'
```

New GET `/ngwaf` /v1/requests endpoint enables searching WAF requests across all workspaces in an account, with each result tagged by workspace_id.

Next-Gen WAF: Added detector_scope to simulator signal responses (2026-06-22)

BLOG ↗ STABLE

Next-Gen WAF Simulate API gains detector_scope field per signal and doubles request body size limit to 200 KB.

└──▷ **TRY IT**

Inspect which detector scope triggered each signal when simulating a reconstructed request — useful for distinguishing built-in WAF rule hits from custom workspace or account rules.

```
$ curl -s -X POST
'https://dashboard.signalsciences.net/api/v0/corps/<corp_name>/sites
/<site_name>/simulate' \
  -H 'x-api-user: <email>' \
  -H 'x-api-token: <token>' \
  -H 'Content-Type: application/json' \
  -d @reconstructed_request.json \
  | jq '.signals[] | {signal: .type, scope: .detector_scope}'
```

Adds `detector_scope` field to every signal in Simulate API responses, indicating whether detection came from a `system`, `workspace`, `account`, or `unknown` scope — matching parity with the `/requests/{id}` API.

Increases Simulate API request body size limit from 100 KB to 200 KB to support reconstructed requests in the import flow.

Fastly Terraform Provider 9.3.0 (2026-06-17) BLOG ↗ STABLE

Fastly Terraform Provider 9.3.0 adds DNS Zones, TSIG Keys, and an immediate-block NGWAF threshold action.

└──▷ **USE IT**

Configure an NGWAF threshold that blocks traffic immediately upon breach rather than waiting for a rate window to close — ideal for high-confidence attack signals.

```hcl
resource "fastly_ngwaf_threshold" "example" {
  # ... existing threshold config ...
  action = "block_immediately"
}
```

Adds Terraform resources for managing DNS Zones and TSIG Keys via the Fastly provider.

Supports `block_immediately` as a new action in NGWAF threshold configurations, enabling instant traffic blocking when a threshold is breached.

Next-Gen WAF: New block_immediately threshold action (2026-06-17) BLOG ↗

STABLE

Next-Gen WAF adds block_immediately threshold action to unconditionally block high-risk signals without a rate window.

└──▷ **HOW TO FIND IT**

Immediately block all requests tagged with the CVE signal at the threshold level, with no rate window required — useful for zero-tolerance enforcement during active exploitation attempts.

📍 In the Next-Gen WAF console, go to Thresholds › select or create a threshold › set Action to block_immediately and choose CVE (or Pentest, XXE, or AWS-SSRF) as the signal.

JavaScript SDK 3.43.0 (2026-06-16) BLOG ↗ STABLE

Fastly JS SDK 3.43.0 adds stale-if-error support for resilient cache serving during origin failures.

Adds stale-if-error cache control support, enabling the SDK to serve stale cached responses when origin errors occur.

Rust SDK 0.13.0 (2026-06-16) BLOG ↗ STABLE

Rust SDK 0.13.0 adds stale-if-error guest caching, async request handoff, and modern Rust lock primitives.

└──▷ **USE IT**

Serve a stale cached response for up to 60 seconds when an origin is unreachable, instead of surfacing a SendError to the client.

```rust
request.set_stale_if_error(Duration::from_secs(60));
let response = match request.send("my_backend") {
    Ok(resp) => resp,
    // SDK automatically substitutes cached stale-if-error response
here
    Err(e) => return Err(e),
};
```

Disable stale-if-error caching entirely so origin errors are always surfaced to the client without serving stale content.

```rust
request.set_stale_if_error(Duration::ZERO);
```

Hand off an in-flight async request so the host delivers the eventual response to the client, freeing your compute to continue other work.

```rust
let pending = request.send_async("my_backend")?;
pending.send_to_client();
```

Adds `PendingRequest::send_to_client` method to hand off responsibility for sending in-flight async responses back to the client to the host.

Supports `stale-if-error` in the HTTP guest cache: serves stale cached responses when `Request::send` or `Request::send_async` would otherwise return a `SendError`.

Enables configurable stale-if-error duration via `Request::set_stale_if_error` and `CandidateResponse::set_stale_if_error`, including opt-out via `Duration::ZERO`.

Exposes `Response::masked_error` on the request collapse leader to surface the underlying error that a stale-if-error response replaced.

Allows 5xx responses to trigger stale-if-error behavior by returning an error from the `after_send` cache hook.

+2 more

WAS THIS USEFUL? 👍 👎

Sublime Security

Feedback ⓘ

Platform v1.85 (ASA Context in User Report Auto-Replies) (2026-07-01) BLOG ↗

Sublime Security is an email security platform that detects and blocks phishing, malware, and other email-based threats using customizable, logic-based detection rules.

Sublime Security Platform v1.85 adds AI-generated explanation variable for user report auto-reply templates.

└──▷ **HOW TO FIND IT**

Automatically tell reporters why a message was flagged, in plain language, by adding the ASA summary to your existing auto-reply template.

📍 In the console, go to Actions › Send User Report Result (or Automations › Close the loop with reporter), open your reply template, and insert {{asa_simple_summary}} where you want the AI explanation to appear.

Adds `{{asa_simple_summary}}` template variable that injects a plain-language AI (ASA) classification explanation into user report auto-reply messages.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Lacework FortiCNAPP

Feedback ⓘ FortiCNAPP Release Updates

2 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Lacework FortiCNAPP is a cloud-native application protection platform that provides security monitoring, threat detection, and compliance management across cloud infrastructure and workloads.

FortiCNAPP adds AI-generated incident reports, FortiGuard IoC feeds, LLM skills scanning, EKS/GKE low-latency audit log alerts, and SAST for Windows across recent releases.

└──▷ TRY IT

Use the AiAssistants API with alert context support to programmatically pull AI-generated triage context for a specific alert.

```
$ curl -X POST 'https://<your-
instance>.lacework.net/api/v2/AiAssistants' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"context": {"alertId": "<alert-id>"}}'
```

Navigate to the Alert Dashboard to triage active threats filtered by severity or sub-category using the new interactive time chart.

📍 In the FortiCNAPP console, go to Alerts › Alert Dashboard, then use the severity, category, or sub-category filters and click a time-chart segment to drill into alerts for that period.

Enable FortiGuard IoC feed integration to enrich detections with Fortinet threat intelligence.

📍 In the FortiCNAPP console, go to Settings › Integrations › Threat Intelligence and enable the FortiGuard IoC Feed integration.

New Alert Dashboard: View and filter alerts by severity, category, or sub-category with an interactive time chart.

Integrates FortiGuard IoC feed as a threat intelligence source for detections.

New Vulnerability Scanner 0.29.0 released with updated scanning capabilities.

Explorer adds natural language search and a redesigned Saved Queries panel.

Cloud activity log pages unified into a single view; AWS CloudTrail and GCP/Azure cloud log pages redesigned.

+28 more

▾ **1 more release in this issue** · seen 2026-07-13

FortiCNAPPRelease Updates NOTES ↗ STABLE

FortiCNAPP adds AI-generated incident reports, FortiGuard IoC feeds, LLM Skills Scanner, EKS/GKE low-latency audit alerting, and SAST for Windows across recent releases.

└──▷ **HOW TO FIND IT**

Investigate a composite alert's AI-generated incident summary and recommended remediation steps without manual triage.

📍 In the FortiCNAPP console, navigate to Alerts › select a composite alert › open the Incident Report or Remediation tab to view AI-generated analysis.

Query alert data using plain English instead of LQL to accelerate threat hunting in Explorer.

📍 In the FortiCNAPP console, go to Explorer › use the natural language search bar to enter a query such as 'show me all failed SSH logins in the last 24 hours', then save results via the redesigned Saved Queries panel.

Retrieve AI-assisted alert context programmatically to enrich SOAR playbooks or ticketing workflows.

```
$ curl -X POST 'https://<instance>.lacework.net/api/v2/AiAssistants'
\
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"type": "AlertContext", "alertId": "<alert_id>"}'
```

New Alert Dashboard lets practitioners view and filter alerts by severity, category, or sub-category.

Integrates FortiGuard IoC feed to enrich threat detection with Fortinet threat intelligence.

New Vulnerability Scanner 0.29.0 released with updated scanning capabilities.

Explorer adds natural language search and a redesigned Saved Queries panel for faster investigation.

Unified view for cloud activity log pages consolidates CloudTrail and other cloud log data.

+38 more

WAS THIS USEFUL? 👍 👎

Aikido Security

Feedback ⓘ

#41 - Aikido MCP, re-testing AI Pentest findings, and Azure Management Group support

13 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Aikido Security is a developer-focused security platform that scans code, dependencies, containers, and cloud infrastructure to detect and prioritize vulnerabilities for fast remediation.

Aikido adds MCP for AI-driven security workflows, AI Pentest re-testing, Azure Management Group auto-discovery, and a wave of new integrations.

└──▷ HOW TO FIND IT

Trigger a targeted re-test of a single AI Pentest finding after applying a fix, instead of re-running the full pentest.

> 📍 In the Aikido console, navigate to AI Pentest › Findings, select the individual finding you fixed, and click Re-test to have Aikido verify the fix in isolation.

Connect an Azure Management Group so Aikido automatically discovers all current and future subscriptions without manual re-configuration.

> 📍 In the Aikido console, go to Integrations › Cloud › Azure, choose Management Group, enter your Management Group ID, and save — Aikido will enumerate all child groups and subscriptions automatically.

New Aikido MCP integration lets AI agents analyze code for security risks and fixes locally, without sending code or data to the cloud.

New re-test capability for individual AI Pentest findings lets teams confirm a specific fix is resolved without running a full pentest.

New Azure Management Group support: connect once and Aikido auto-discovers all existing and future groups and subscriptions.

New Broker Helm chart released for self-hosted broker deployments.

New GitHub Enterprise Server on-premises support via Aikido Broker for code scanning.

+19 more

▾ 12 more releases in this issue · seen 2026-07-13

> #42 - Dep Upgrade Analysis, Eclipse IDE, AI Pentest Preview & Real Time Kubernetes Ima
>
> NOTES ↗ STABLE

Aikido adds Kubernetes real-time image scanning, Eclipse IDE plugin, AI pentest preview, and dependency upgrade breaking-change analysis.

New Dependency Upgrade Breaking Change Analysis: Aikido AutoFix surfaces code-breaking impacts of dependency upgrades before they are applied.

New Kubernetes In-Cluster Image Scanning: continuously monitors container images running in live Kubernetes clusters for vulnerabilities.

New AI Pentest Skip-Payment model: run a full whitebox pentest and pay only to unlock details of High or Critical findings.

New Eclipse IDE plugin: scans projects in real time for secrets, API keys, and security vulnerabilities as you code.

New and improved Reachability view: refreshed design with deeper, more granular insights accessible directly from the Feed on hover.

+19 more

#43 - Aikido Infinite (Continuous Pentesting), VM Reachability Diagrams & IDE full sca

NOTES ↗ STABLE

Aikido adds continuous pentesting, VM reachability diagrams, full IDE workspace scanning, and expanded integrations.

└──▷ **HOW TO FIND IT**

Visualize the network path to a vulnerable cloud VM to quickly assess real-world exploitability before prioritizing remediation.

> 📍 In the Aikido console, open a Cloud VM finding, then select Reachability Diagram to view the network path, entry points, and exposed ports leading to the affected resource.

Attach custom Jira fields (e.g. team or severity tier) when Aikido creates issues, so tickets route correctly without manual editing.

> 📍 In the Aikido console, go to Integrations › Jira, open your integration settings, and configure Custom Fields with the desired Jira field keys and values.

New Aikido Infinite (Continuous Pentesting): automatically re-tests your app after every code change, focusing on what changed and validating exploitability to reduce false positives.

New Reachability Diagrams for Cloud VMs: visualizes whether a finding is exploitable by mapping the network path, entry points, and exposed ports in your cloud environment.

Full Workspace Scan now available across all supported IDEs (VS Code, JetBrains, Visual Studio, Eclipse, Cursor), with IaC scanning added for VS Code.

Adds package health scoring for PHP and malware detection for skills.sh, Go, Rust, and WordPress plugins.

Container image scanning supports include/exclude filtering by pattern.

+8 more

#44 - Talk to Aikido Autofix, Pentest Threat Model and support for Alibaba Cloud

NOTES ↗ STABLE

Aikido adds conversational AutoFix refinement, Pentest Threat Model, Alibaba Cloud support, and new APIs in release #44.

└──▷ HOW TO FIND IT

Iteratively refine an AutoFix to match your team's code conventions using natural language inside Aikido.

> 📍 In the Aikido UI, open an AutoFix suggestion, click `Refine with Aikido AI`, and type a natural-language instruction such as 'Use async/await instead of callbacks and follow our error-handling conventions'.

Adds `Refine with Aikido AI` — iteratively adjust AutoFixes using natural language to match your code style without leaving Aikido.

New Pentest Threat Model provides a real-time, transparent view of your attack surface and how your application's security posture will be tested.

Supports Alibaba Cloud: scan for misconfigurations, exposed resources, and risky settings across your Alibaba Cloud environment.

Adds custom context on SAST rules to improve AutoTriage accuracy.

Displays CVEs on VM attack path diagrams in Cloud & VM views.

+9 more

#45 - Aikido × Lovable - Vibe, Fix, Ship and Container Reachability NOTES ↗

STABLE

Aikido adds Lovable pentesting integration, container network reachability, new APIs, and expanded Zen Firewall agent support.

Adds Container Reachability: maps the complete network route to running containers and surfaces CVEs in deployment context, showing what is exposed to the internet.

Detects `org.opencontainers.image.source` tags in Docker images for improved

container identification.

Extends Virtual Machine Reachability support to Azure.

Refreshes Zen Firewall installation flow with new installation instructions.

Ships Zen Firewall agents for .NET, PHP, Ruby, Golang, and Node.

+5 more

#46 - Escalating your AI Pentest and Additional Context for AutoTriage

NOTES ↗ STABLE

Aikido adds AI pentest escalation, AutoTriage custom context, cloud reachability for GCP, and new Zen firewall agents.

Enables escalation of AI pentest vulnerabilities beyond default cautious mode to perform deeper exploitation analysis and confirm whether protection layers hold up.

Adds custom environment context input for AutoTriage to improve finding accuracy and reduce false positives.

New out of SLA graph added to Reports Trends Over Time view.

Supports direct reply or ignore actions via GitLab MR comments integration.

Extends cloud reachability analysis for containers to Google Cloud Run.

+9 more

#47 - Aikido Device Protection and Aikido AI Chat NOTES ↗ STABLE

Aikido adds Device Protection for blocking malicious packages on developer machines, AI Chat in AutoFix, and new API endpoints.

└──▷ **HOW TO FIND IT**

Use the AI Chat inside the AutoFix modal to refine a generated fix without leaving the product — useful when the suggested PR needs a small scope adjustment.

📍 In the console, open a finding in the Code Analysis or AutoFix modal, then use the embedded Aikido AI Chat panel to describe the change you want (e.g., 'Scope the fix to only the affected function, not the entire file') and submit.

Add organization-specific context (custom libraries, internal policies) so AutoFix generates more accurate PRs tailored to your codebase.

```
$ # Place a .aikido file at the root of your repository with custom
context
echo 'We use our internal wrapper `safe_query()` instead of raw DB
calls.' > .aikido
```

New Device Protection agent blocks malicious packages (npm, PyPI, Maven, NuGet, and more) and risky IDE/browser extensions at install time on developer devices to stop supply chain attacks.

Aikido AI Chat is now embedded in the AutoFix & Code Analysis Modal, enabling natural-language interaction with security findings and in-product AutoFix PR customization.

Adds cloud support for detecting Supabase common configuration issues.

Enables custom context for AutoFix via `.aikido` files.

Aikido Infinite now supports custom scan frequencies (e.g., every 6 weeks).

+5 more

#48 - AI Code Audit, Aikido MCP and ServiceNow Integration NOTES ↗ STABLE

Aikido adds AI code auditing, MCP integration, ServiceNow ticketing, Oracle/Alibaba cloud support, and Zen Firewall agent expansion.

└──▷ **HOW TO FIND IT**

Query your Aikido security findings — such as leaked secrets or critical open vulnerabilities — directly from an AI coding assistant via the MCP integration.

> 📍 In your AI coding assistant's MCP settings, add the Aikido MCP server, then prompt: "Show me all critical findings and leaked secrets in the <repo> repository."

Automatically route Aikido security findings into ServiceNow for teams that manage remediation through existing ITSM workflows.

> 📍 In the Aikido console, go to Integrations › Task Trackers › ServiceNow and connect your ServiceNow instance to enable automatic ticket creation from findings.

Audit source code for logic flaws and vulnerabilities beyond what traditional SAST tools detect, using AI-powered pentest-grade analysis.

> 📍 In the Aikido console, go to Code › AI Code Audit and connect a repository to trigger pentest-grade AI reasoning on your source code.

New AI Code Audit capability performs pentest-grade reasoning on source code to catch logic flaws and vulnerabilities that traditional scanners miss — activated by connecting a repo.

New Aikido MCP integration lets AI coding assistants (e.g. Cursor, Claude) query Aikido issues — critical findings, leaked secrets, and open vulnerabilities — directly from the coding environment.

New ServiceNow integration automatically converts security findings into ServiceNow tickets, routing them into existing team workflows.

Adds CVE lookup via Aikido's chat interface for rapid vulnerability research.

Adds release gating with SLA Mode for code scanning, enabling enforcement of security SLAs at the pipeline level.

+11 more

#49 - Visualize Threat models, Device Protection for Windows and Labels on Repositorie

NOTES ↗ STABLE

Aikido #49 adds threat model diagrams, Windows device protection, repo/container labels, and a wave of new API endpoints.

└──▷ HOW TO FIND IT

Filter your issue feed to only repositories tagged with a specific business unit or data-sensitivity label to prioritize remediation for PII-handling services.

📍 In the Aikido console, go to Repositories (or Container Images) › select a repo › Labels › add a label (e.g. `PII` or `business-unit:payments`). Then in the Issues feed, use the Labels filter to scope results to that label.

Export an SBOM for a specific device to inventory software components and feed downstream vulnerability management tools.

📍 In the Aikido console, go to Endpoint › Devices › select a device › Export SBOM.

New threat model visualization: every AI pentest and code audit now ships an architecture diagram showing roles and trust boundaries.

Device Protection now supports Windows, including MDM rollout for fleet-wide deployment.

New Labels feature lets you attach arbitrary metadata (business unit, PII, region, etc.) to repositories and container images and filter your issue feed by those labels.

Code Audit gains a Retest capability to re-evaluate issues after code changes.

Endpoint feature adds device SBOM export.

+7 more

#50 - Agentic Dependency AutoFix, Registry Proxy and Device Protection adds Ruby and R

NOTES ↗ STABLE

Aikido adds agentic dependency AutoFix, a registry proxy for malware blocking, and Ruby/Rust device protection.

└──▷ **HOW TO FIND IT**

Use Agentic Dependency AutoFix on a project with many interlocking dependencies to let Aikido automatically compute and apply a safe upgrade path.

> 📍 In the Aikido console, open a dependency vulnerability finding and select `AutoFix` to trigger the agentic resolution workflow.

New Agentic Dependency AutoFix automatically resolves complex, interlocking dependency interactions on your behalf.

New Registry Proxy routes your private registry through Aikido to block malware and risky new packages and serve safe, patched versions before they reach your codebase.

Device Protection now covers Ruby (gems) and Rust (crates), blocking malware on install and adding them to your device SBOM.

Continuous Pentest adds Azure support for cloud environment scanning.

Continuous Pentest now supports selecting specific branches for `On Schedule` runs.

`+8 more`

`New - Code Audit: Pentest-grade reasoning on your source code, no environment required`

NOTES ↗ **STABLE**

Aikido launches Code Audit: agentic pentest-grade source code reasoning that finds logic flaws without a live environment.

└──▷ **HOW TO FIND IT**

Launch a code audit on a connected repo directly from the Aikido console — no staging URL or pentest scope needed.

> 📍 In the Aikido console, go to Code Audit, select the target codebase, review the credit cost shown upfront, and click Start Audit.

New Code Audit feature applies agentic, pentest-grade reasoning directly to source code to detect logic flaws — such as cross-tenant data leakage — that pattern-based scanners miss.

Supports codebases that are difficult to test via URL: mobile apps, smart contracts, desktop apps, embedded software, and libraries.

Integrates with Aikido's existing issue experience: findings include code-based evidence, root cause analysis, remediation guidance, and AI-powered AutoFix for opening remediation PRs.

Upfront credit-based pricing displayed before audit starts, with cost scaled to codebase size and complexity — no hidden setup costs.

Requires no staging environment, traffic replay, or agent deployment — connect a repo and start the audit immediately.

New - Extreme Noise Reduction for Secrets and CVEs 🚀 NOTES ↗ STABLE

Aikido adds LLM-powered AutoTriage for secrets and CVEs to cut noise and automate prioritization decisions.

└──▷ HOW TO FIND IT

Review a Secrets AutoTriage outcome to confirm or override the agent's liveness and scope assessment for an ambiguous credential.

📍 In the Aikido console, open a secrets finding › AutoTriage result to read the agent's explanation of the credential type, liveness, and risk level, then accept or override the decision.

New Secrets AutoTriage agent: assesses whether a detected secret is a real, exposed credential and evaluates its liveness and scope to determine actionable risk.

New CVE Exploitability Analysis agent: uses LLM reasoning against your actual codebase to determine whether a CVE is reachable and how urgently it must be addressed.

Automated remediation actions: the CVE agent can upgrade, downgrade, snooze, or ignore findings automatically based on configurable rules.

Audit log per run: every agent analysis is logged with a short rationale so teams can review decisions and override them.

New - Protect Company Devices From Supply Chain Attacks 🚀 NOTES ↗ STABLE

Aikido launches Device Protection: real-time malware blocking, extension governance, and group-based policies for developer devices.

└──▷ HOW TO FIND IT

Review and action a developer's blocked-install exception request without leaving Aikido.

📍 In the Aikido console, go to Device Protection › Inbox, locate the pending approval request, review the flagged package details, and approve or deny the exception.

New Device Protection agent provides real-time malware blocking across npm, PyPI, Maven, and NuGet package ecosystems powered by Aikido Intel.

Adds extension governance covering VS Code, OpenVSX, Chrome, and skills.sh extensions with policy-based controls.

New request-and-approval workflow pauses blocked installs and routes one-click exception requests to the admin Inbox inside Aikido.

> Supports group-based policies allowing different installation rules to be set per team.
>
> Continuous device inventory delivers an always-current view of every package and extension installed per developer device.

WAS THIS USEFUL? 👍 👎

Gruntwork cloud-nuke

Feedback ⓘ v0.52.0 NOTES ↗

CODE ↗

A tool for cleaning up your cloud accounts by nuking (deleting) all resources within it

cloud-nuke v0.52.0 adds parallel scanning and deletion via a new `--parallelism` flag, dramatically speeding up multi-region cleanups.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.52.0 https://github.com/gruntwork-io/cloud-nuke.git
# already have the repo? check out this version:
$ git checkout v0.52.0
```

└──▷ TRY IT

Speed up a full AWS account cleanup across many regions by raising the concurrency well above the default of 10.

```
$ cloud-nuke aws --region us-east-1 --region us-west-2 --region eu-west-1 --parallelism 25
```

Throttle parallelism to 3 when nuking a large account where API rate limits are a concern, while still scanning concurrently.

```
$ cloud-nuke aws --parallelism 3 --resource-type ec2 --resource-type s3 --dry-run
```

Adds `--parallelism` flag (default 10) to control concurrent scanning and deletion of regions and resource types under a single global concurrency cap.

Enables concurrent nuking of regional resources while preserving within-region dependency order; global resources (IAM, Route53, CloudFront) are deleted last.

Breaking change for library consumers: `aws.NukeAllResources` and `gcp.NukeAllResources` now require an additional `parallelism int` parameter.

└──▷ BREAKING ON UPGRADE

! Library consumers: `aws.NukeAllResources` and `gcp.NukeAllResources` now take an additional `parallelism int` parameter.

WAS THIS USEFUL?  

Sysdig Secure

Feedback ⓘ Secure SaaS - April 15, 2026

9 RELEASES BLOG ↗ LATEST STABLE

Sysdig Secure is a cloud-native security platform that detects threats, manages vulnerabilities, ensures compliance, and secures containers, Kubernetes, and cloud environments at runtime.

Agentless Host Scanning now covers Windows, Fedora, and AWS Bottlerocket hosts via snapshot-based scanning.

Adds agentless snapshot-based host vulnerability scanning for Windows hosts.

Adds agentless snapshot-based host and container vulnerability scanning for Fedora hosts.

Adds agentless snapshot-based host and container vulnerability scanning for AWS Bottlerocket hosts.

▾ 8 more releases in this issue

> Secure SaaS - July 1, 2026 BLOG ↗ STABLE
>
> Registry Scanner v0.12.0 adds OCI image format and multi-architecture bundle scanning across all supported registries.
>
> Supports scanning of OCI image manifests (`application/vnd.oci.image.manifest.v1+json`), covering single-platform images published in OCI format.
>
> Supports OCI image indexes and Docker manifest lists (`application/vnd.oci.image.index.v1+json` and `application/vnd.docker.distribution.manifest.list.v2+json`), resolving multi-architecture images to a platform-specific manifest for scanning.
>
> Recognizes image manifests that omit the optional `mediaType` field (e.g., those produced by buildah), which were previously skipped as unsupported.
>
> Adds intelligent multi-architecture platform selection: prefers `linux/amd64`, then any `linux/*`, then the first scannable manifest; skips attestation manifests and `unknown/unknown` entries.
>
> Publishes the Registry Scanner job image as a multi-architecture manifest (`linux/amd64` and `linux/arm64`), enabling native support for ARM-based clusters.

> Secure SaaS - July 9, 2026 BLOG ↗ STABLE
>
> Host Scanner adds Google COS/GKE node vulnerability coverage and improved Ruby gem file-level attribution.

Supports Google Container-Optimized OS (COS): Host Scanner now identifies installed packages on COS hosts (e.g., GKE nodes), extending vulnerability scanning to those workloads.

Improves Ruby package detection: Host Scanner tracks source `.rb` files within installed gems, enabling more precise file-level vulnerability attribution.

Secure SaaS - June 25, 2026 BLOG ↗ STABLE

Sysdig CLI Scanner v1.27.2 adds detection support for Google Container-Optimized OS (COS) packages.

Secure SaaS - June 26, 2026 BLOG ↗ STABLE

Sysdig Secure refreshes the Threats UI with bulk actions and resource focus, and adds Azure Monitor Logs Ingestion API for Sentinel.

└──▷ **HOW TO FIND IT**

Switch your Sentinel forwarding integration to the new Azure Monitor Logs Ingestion API before Microsoft's HTTP Data Collector API reaches end of life.

📍 In Sysdig Secure, go to Integrations › Forwarding, select your Microsoft Sentinel destination, and choose the Azure Monitor Logs Ingestion API authentication method, then save.

Adds bulk actioning on Threat aggregations, allowing review of multiple occurrences at once from a single pane.

Threats now scoped to a single resource, surfacing context-aware risk signals per asset.

New occurrence tabs let practitioners toggle between the latest and prior threat occurrences within the same view.

Supports the Azure Monitor Logs Ingestion API for Microsoft Sentinel SIEM integration, replacing the deprecated HTTP Data Collector API.

Secure SaaS - June 9, 2026 BLOG ↗ STABLE

Sysdig Secure adds end-of-life visibility and policy enforcement for OS and Go runtime components in Vulnerability Management.

└──▷ **HOW TO FIND IT**

Enforce a pipeline gate that fails builds when any scanned component is end-of-life or within 90 days of EOL — configured via a Component Lifecycle rule in a Rule Bundle.

📍 In the Sysdig Secure console, go to Policies › Rule Bundles, create or edit a bundle, add a new rule of type `Component Lifecycle`, and set the condition to fail when a

component has reached its end-of-life date or is within 90 days of it. Attach the bundle to the relevant pipeline policy.

Consume lifecycle metadata from the CLI scanner JSON output in a CI pipeline to programmatically detect EOL components without a UI.

```
$ sysdig-cli-scanner <image> --json-scan-result result.json && jq '.info.resultDetails.packages[] | select(.lifecycle.status == "end-of-life")' result.json
```

Adds end-of-life (EOL) visibility for operating systems and the Go runtime in Vulnerability Management, surfacing lifecycle status alongside vulnerability data in the Resource Details drawer, Components tab, Findings Details drawer, CVE Details drawer, and scan result side panels.

New Lifecycle column and filter in the Components tab lets practitioners quickly identify components no longer receiving vendor support.

CLI scanner JSON output and the scan results API now include lifecycle metadata, enabling pipeline and integration consumers to act on EOL status programmatically.

New Component Lifecycle rule type in Rule Bundles enables policy evaluations to fail when a component has reached its EOL date or is within 90 days of it.

Sysdig CLI Scanner v1.27.1 adds end-of-life support for Vulnerability Management.

Secure SaaS - May 15, 2026 BLOG ↗ STABLE

Sysdig Secure adds Response Actions directly in the resource drawer for hosts, Kubernetes workloads, and AWS resources.

└──▷ HOW TO FIND IT

Contain an active incident on a Kubernetes workload without leaving the resource view — inspect the workload in the resource drawer, then trigger a Response Action inline.

📍 In Sysdig Secure, navigate to the resource drawer for a Kubernetes workload › select the Response Actions tab › choose a compatible Kubernetes action and execute it.

Audit past containment decisions on an AWS resource by reviewing execution status and configuration history in one place.

📍 In Sysdig Secure, open the resource drawer for an AWS resource › select the Response History tab › inspect each action's execution status, target information, and configuration.

Enables Response Actions (threat hunting, incident containment) directly from the resource drawer without navigating away from the resource view.

Supports host, container, and file Response Actions for hosts when Host Shield is configured.

Supports Kubernetes Response Actions for workloads when Cluster Shield is enabled.

Supports Cloud Response Actions for AWS resources when Cloud Response Actions are deployed.

Adds a Response History panel in the resource drawer to review past actions, execution status, target info, and configuration.

Secure SaaS - May 26, 2026 BLOG ↗ STABLE

Sysdig CLI Scanner v1.27.0 adds Flashpoint KEV tag and Flashpoint CVSSv3 Temporal score to vulnerability output.

Adds Flashpoint KEV (Known Exploited Vulnerabilities) tag to CLI Scanner vulnerability output for faster exploit-aware triage.

Adds Flashpoint CVSSv3 Temporal score to CLI Scanner vulnerability output, surfacing real-world exploitability context alongside base scores.

Secure SaaS - May 5, 2026 BLOG ↗ STABLE

Sysdig Secure adds Response History page for auditing all Response Actions across events, automations, and APIs.

└──▷ **HOW TO FIND IT**

Review and roll back a containment action that was triggered incorrectly by an automation rule.

📍 In Sysdig Secure, go to Response History, locate the containment action by timestamp or actor, review the collected artifacts, and select Revert to undo the action.

New Response History page centralizes all Response Actions taken manually, via Automations, or through APIs — showing what was done, where, when, and by whom.

Enables retrieval of collected artifacts and reversal of containment actions performed in error from a single audit view.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Tencent Cloud CubeSandbox

Feedback ⓘ v0.5.1 NOTES ↗

Instant, Concurrent, Secure & Lightweight Sandbox for AI Agents.

CubeSandbox v0.5.1 adds a standalone lifecycle manager, three-value timeout semantics, host-mount allowlisting, and a Node.js SDK preview.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.5.1
https://github.com/TencentCloud/CubeSandbox.git
# already have the repo? check out this version:
$ git checkout v0.5.1
```

└──▷ USE IT

Pin a long-running sandbox so it is never reclaimed by the idle sweeper — useful for persistent dev environments or always-on agent workloads.

python

```
# Python SDK
sandbox = await
client.create(options=CreateOptions(timeout=NEVER_TIMEOUT))

# Adjust an already-running sandbox
await sandbox.set_timeout(NEVER_TIMEOUT)
```

Restrict which host directories sandboxes may bind-mount to prevent arbitrary host path exposure in a multi-tenant cluster.

yaml

```
# cubemaster config (yaml)
host_mount:
  allowed_prefixes:
    - /data/shared/
    - /data/project-a/
  # '/' is explicitly forbidden and will be rejected at admission
```

Scale CubeProxy horizontally with CLM coordinating auto-pause across replicas, configured via Terraform.

```hcl
# terraform.tfvars
TENcentCLOUD_CUBE_PROXY_REPLICAS = 3
cube_lifecycle_manager_replicas  = 1
cube_lifecycle_manager_discovery_refresh_interval = 30
```

Introduces cube-lifecycle-manager (CLM), a standalone control-plane service extracted from the CubeProxy sidecar, enabling CubeProxy to scale horizontally with Redis-based replica discovery and coordinated auto-pause/resume.

Adds three-value timeout semantics: omitted inherits cluster default, `NEVER_TIMEOUT` (-1) disables idle reclaim, `0` triggers immediate reclaim, and N>0 sets an idle TTL in seconds — with server-side defaults in CubeMaster.

Adds set_timeout() / SetTimeout() to Python and Go SDKs (including `NEVER_TIMEOUT` sentinel) and exposes timeout on sandbox create in the Web UI.

Adds host-mount path allowlisting: bind-mounts are restricted to configurable prefixes (default `/data/shared/`) with `filepath.Clean` traversal neutralization; `/` is explicitly forbidden.

Adds Node.js / TypeScript SDK preview (`@cubesandbox/sdk`) aligned with the Python and Go surface.

+7 more

└─▷ BREAKING ON UPGRADE

- ! Host-mount paths are now restricted to configurable prefixes (default `/data/shared/`); previously any absolute path was accepted, so existing sandboxes using host bind-mounts outside the allowed prefix will be rejected.
- ! Go and Python SDK `CreateOptions.Timeout` (and equivalent fields) are now optional pointers; code that relied on the SDK injecting a hard-coded default timeout will no longer receive that default — the cluster `default_timeout_insec` is used instead, or never-timeout if unset.
- ! The `cube-proxy-sidecar` in-image AutoPause coordinator has been extracted into a standalone `cube-lifecycle-manager` service; deployments that relied on the sidecar for AutoPause coordination must now deploy the separate CLM service.
- ! CubeEgress transparent-proxy listen address is no longer hard-coded to `192.168.0.1`; it is now derived from `CUBE_SANDBOX_NETWORK_CIDR`, so deployments using custom sandbox network CIDRs will get a different TPROXY/OpenResty listen address.
- ! Go SDK `envd` RPCs (commands, files, filesystem, PTY) now route to port 49983 instead of port 49999; integrations that hard-coded port 49999 for these operations will break.

WAS THIS USEFUL? 👍 👎

NoLabs nono

Feedback ⓘ v0.67.0

5 RELEASES · 2026-06-15 → 2026-07-06 NOTES ↗ CODE ↗ LATEST STABLE

Sandbox any AI agent in seconds - zero setup, zero latency.

nono v0.67.0 adds a standalone proxy command, AWS SigV4 MitM auth, declarative OAuth capture, resource limiting, exec intercept, and richer sandbox token expansion.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.67.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.67.0
```

└──▷ TRY IT

Run the credential/filtering proxy on a fixed port with a specific domain allowlist — useful when an agent or CI job needs a pre-started proxy to route through.

```
$ nono proxy --port 8080 --allow-domain api.github.com
```

Extend a community registry profile locally so you can add custom rules without forking the upstream profile.

```
$ nono profile init opencode --extends always-further/opencode
```

Start a loopback proxy with no token auth for local development or testing pipelines that don't require authentication.

```
$ nono proxy --port 8080 --no-auth
```

Adds `nono proxy` standalone command to run the network-filtering and credential-injection proxy as an independent server, with options for fixed port, no-auth mode, and domain allowlisting.

Implements AWS SigV4 authentication for the MitM proxy, enabling signed AWS API request interception and credential injection.

Adds declarative sandboxed OAuth capture, allowing profiles to define OAuth flows without manual credential handling.

Introduces resource limiting support in sandbox profiles.

Adds `exec` intercept action for tool-sandbox command policies, enabling a helper binary to replace a matched command's real binary inside the existing sandbox.

+5 more

▾ **4 more releases in this issue** · 2026-06-15 → 2026-07-06

v0.66.0 NOTES ↗ CODE ↗ STABLE

nono v0.66.0 adds HTTP/2 proxy support, a @git:common-dir dynamic token, and a simplified self-invocation policy for tool sandboxes.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.66.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.66.0
```

└──▷ **TRY IT**

Enable HTTP/2 multiplexing when running an agent that makes many concurrent requests through the credential proxy (e.g., Maven/Gradle artifact downloads) to improve throughput.

```
$ nono run --profile always-further/opencode --allow-http2 --
opencode
```

Adds `--allow-http2` flag to enable HTTP/2 multiplexing (via ALPN) for reverse proxy and credential injection upstream connections.

Adds `@git:common-dir` dynamic token for tool-sandbox profiles, allowing policies to reference the Git common directory at runtime.

Simplifies self-invocation policy for sandboxed tools, reducing boilerplate needed when a tool is permitted to call itself.

v0.65.0 NOTES ↗ CODE ↗ STABLE

nono v0.65.0 introduces tool sandboxing, giving delegated tools their own isolated child sandboxes with separate filesystem, network, and credential policies.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.65.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.65.0
```

└──▷ **TRY IT**

Restrict `gh` so the agent can only list and view issues — not comment — and inject a credential-proxied GitHub token scoped to those API paths only.

```
$ nono run --profile opencode -- opencode
```

Diagnose why the tool sandbox blocked a `gh issue comment` invocation — useful when tuning `command_policies` argv allow-lists.

```
$ nono why --profile gh --command gh -- issue comment 1052
```

Define a tool sandbox policy in your profile that scopes `gh` to read-only issue access with a proxied credential and an L7 endpoint allow-list.

```json
{
  "command_policies": {
    "credentials": {
      "github-api": {
        "type": "proxy",
        "upstream": "https://api.github.com",
        "credential_key": "keyring://gh:github.com/example?
decode=go-keyring",
        "env_var": "GH_TOKEN",
        "inject_header": "Authorization",
        "credential_format": "Bearer {}"
      }
    },
    "commands": {
      "gh": {
        "from": {
          "session": {
            "sandbox": {
              "fs_read": ["."],
              "credentials": [
                {
                  "name": "github-api",
                  "endpoint_policy": {
                    "default": "deny",
                    "allow": [
                      { "method": "GET", "path": "/repos/nolabs-
ai/nono/issues/**" }
                    ]
                  }
                }
              ]
            },
            "invocation_policy": {
              "default": "deny",
              "allow": [
                { "argv": { "prefix": ["issue", "list"] } },
                { "argv": { "prefix": ["issue", "view"] } }
              ]
            }
          }
        }
      }
    }
  }
}
```

Adds tool sandbox (ETI): delegated tools such as `git`, `gh`, and `curl` now run in isolated child sandboxes with their own `fs_read` / `fs_write`, network rules, credential grants, and argv invocation policies — independent of the agent's session sandbox.

Enables per-tool credential proxy injection via `command_policies.credentials`, letting profiles expose a proxied token to a specific tool (e.g. `gh`) while enforcing L7 endpoint and HTTP-method allow-lists on that token.

Extends `nono why` to diagnose tool-sandbox argv policy denials via `--command <cmd> -- <args…>`, in addition to its existing filesystem and network queries.

└──▷ **BREAKING ON UPGRADE**

! Custom credentials now require explicit activation; previously working custom credential configurations will no longer function without being explicitly activated.

v0.64.0 NOTES ↗ CODE ↗ STABLE

nono v0.64.0 adds AWS auth config, structured diagnostics JSON output, CI environment detection, and macOS blocked-grants visibility.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.64.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.64.0
```

└──▷ **TRY IT**

Capture machine-readable session diagnostics after a run to feed into a SIEM or audit pipeline.

```
$ nono run --profile always-further/opencode --diagnostics-json --
opencode 2>session-diagnostics.json
```

Adds `aws_auth` config support, enabling native AWS credential handling in agent sandbox profiles.

New `--diagnostics-json` flag on `nono run` emits structured session diagnostics as JSON on stderr, including proxy diagnostics when present.

Update-check now detects CI environments, suppressing spurious upgrade prompts in automated pipelines.

Capability summary now surfaces blocked macOS grants, giving practitioners a clearer view of what the sandbox is denying on macOS.

Exposes structured diagnostics for library and FFI clients, enabling programmatic access to session diagnostic data from Rust, Python, TypeScript, and Go bindings.

v0.63.0 NOTES ↗ CODE ↗ STABLE

nono v0.63.0 adds static env injection, keyring timeout control, $`PACK_DIR` in session hooks, and XDG state dir support.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.63.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.63.0
```

└──▷ **TRY IT**

Inject a fixed environment variable into every sandboxed session — useful for pinning model names or feature flags without leaking them through the host environment.

```
$ nono run --profile always-further/opencode -- opencode
```

Prevent indefinite hangs when the system keyring is locked or unavailable — useful in headless CI environments where keychain prompts never appear.

```
$ NONO_KEYRING_TIMEOUT_SECS=5 nono run --profile always-
further/opencode -- opencode
```

Adds `set_vars` profile field for static environment variable injection into sandboxed sessions — set fixed env values without exposing them through the agent's environment.

Adds `NONO_KEYRING_TIMEOUT_SECS` environment variable to control how long nono waits for keychain/keyring access before timing out.

Adds `$PACK_DIR` expansion support in `session_hooks`, enabling hooks to reference the store pack directory for pack-aware scripts.

Moves CLI runtime state to XDG state directories, aligning with the XDG Base Directory specification for cleaner system integration.

└──▷ **BREAKING ON UPGRADE**

! Runtime state is moved to XDG state directories (`feat(cli): move runtime state to XDG state dirs`), which may break setups that reference the previous runtime state paths.

WAS THIS USEFUL? 👍 👎

NVIDIA OpenShell

Feedback ⓘ v0.0.76

4 RELEASES · 2026-06-29 → 2026-07-03 NOTES ↗ LATEST STABLE

OpenShell is the safe, private runtime for autonomous AI agents.

OpenShell v0.0.76 adds SELinux bind-mount labels and numeric UID support in sandbox policies.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.76 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.76
```

Adds SELinux label support for bind mounts in Docker and Podman container runtimes.

Enables numeric UIDs (in addition to named users) when specifying sandbox process identity in policies.

▾ 3 more releases in this issue · 2026-06-29 → 2026-07-03

v0.0.75 NOTES ↗ STABLE

OpenShell v0.0.75 adds TLS termination for Envoy Gateway ingress and a new manifest-driven gator agent.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.75 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.75
```

Adds TLS termination support for Envoy Gateway ingress via Helm, enabling encrypted traffic handling at the ingress layer.

Introduces a manifest-driven gator agent, allowing agent behavior to be defined and driven through declarative manifests.

v0.0.74 NOTES ↗ STABLE

OpenShell v0.0.74 adds a combined topology config surface for Kubernetes deployments.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.74 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.74
```

Adds combined topology config surface for Kubernetes, enabling unified topology configuration in a single surface.

Breaking: removes `SandboxTemplate.volume_claim_templates` from the API — configurations referencing this field must be updated.

└──▷ BREAKING ON UPGRADE

! The `SandboxTemplate.volume_claim_templates` field has been removed from the API

v0.0.72 NOTES ↗ STABLE

OpenShell v0.0.72 adds AWS SigV4 proxy signing, JSON-RPC/MCP policy enforcement, and Kubernetes agent-sandbox v1beta1 support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.72 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.72
```

Adds proxy-side AWS SigV4 credential signing for CONNECT tunnels in sandbox environments.

Adds L7 policy enforcement for JSON-RPC and MCP traffic.

Supports Kubernetes agent-sandbox v1beta1 API.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Open Policy Agent Gatekeeper

Feedback ⓘ v3.23.0

NOTES ↗

Policy Controller for Kubernetes

Gatekeeper v3.23.0 adds structured log fields, remote cluster status routing, and operation-scoped mutation rules.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.23.0 https://github.com/open-policy-agent/gatekeeper.git
# already have the repo? check out this version:
$ git checkout v3.23.0
```

└──▷ USE IT

Scope a mutation to fire only on CREATE and UPDATE operations, preventing it from running on DELETE — useful when a mutator sets defaults that are irrelevant during deletion.

yaml

```yaml
apiVersion: mutations.gatekeeper.sh/v1alpha1
kind: Assign
metadata:
  name: set-default-label
spec:
  applyTo:
  - groups: [""]
    kinds: ["Pod"]
    versions: ["v1"]
    operations: ["CREATE", "UPDATE"]
  location: "metadata.labels.env"
  parameters:
    assign:
      value: "production"
```

Adds structured field markers to semantic log lines for easier machine-parseable log filtering and SIEM ingestion.

Implements status resource routing for remote cluster mode, enabling multi-cluster Gatekeeper deployments to correctly surface constraint status from remote clusters.

Adds `operations` field to mutation `ApplyTo` spec, allowing mutations to be scoped

to specific admission operations (e.g., CREATE, UPDATE, DELETE).

WAS THIS USEFUL? 👍 👎

Inspektor Gadget

Feedback ⓘ v0.54.0 NOTES ↗ CODE ↗

Inspektor Gadget is a set of tools and framework for data collection and system inspection on Kubernetes clusters and Linux hosts using eBPF

Inspektor Gadget v0.54.0 adds Cilium NetworkPolicy generation, Helm resource limits, and improved exec arg capture.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.54.0 https://github.com/inspektor-gadget/inspektor-gadget.git
# already have the repo? check out this version:
$ git checkout v0.54.0
```

└──▷ **TRY IT**

Generate a Cilium-native NetworkPolicy by observing live traffic, ready to apply directly to a Cilium-managed cluster.

```
$ kubectl gadget run advise_networkpolicy:latest --policy-format cilium
```

Constrain Inspektor Gadget DaemonSet resource usage at deploy time to satisfy admission policies or resource quotas.

```
$ helm upgrade --install gadget gadget/gadget --set resources.requests.cpu=100m --set resources.requests.memory=128Mi --set resources.limits.cpu=500m --set resources.limits.memory=512Mi
```

Adds `--policy-format cilium` option to `advise_networkpolicy`, enabling generation of Cilium `NetworkPolicy` resources in addition to the existing Kubernetes `NetworkPolicy` output.

Enables CPU and memory `requests` and `limits` for the Inspektor Gadget DaemonSet via the Helm chart `resources` field, so deployments fit within cluster resource quotas.

Improves `trace_exec` to read process arguments from the new process address space, capturing more accurate argv data for executed processes.

Adds support for profile names in OpenTelemetry profiles (`otel-profiles`).

Introduces a user-provided BTF cache in the eBPF subsystem (ebpf v0.22.0 upgrade), reducing repeated BTF processing overhead at scale.

└─▷ **BREAKING ON UPGRADE**

! SBOMs are removed from the image index and signed separately, which may break workflows that retrieve SBOMs directly from the image index.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Hunters Pathfinder AI

Feedback ⓘ snapshot-20260713

NOTES ↗

Hunters Pathfinder AI is an AI-powered security assistant that helps analysts investigate, triage, and respond to threats within the Hunters SOC platform.

Hunters adds unified anomalous login detector, Windows password-change enrichment/scoring, and four new AXON threat hunt reports.

└──▷ HOW TO FIND IT

Review AXON proactive threat hunt findings for your environment after a new report campaign is published.

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

New detector: Context Unified Anomalous Login detects successful anomalous logins across identity providers, SaaS apps, cloud platforms, VPNs, and other integrated login sources using per-user behavioral baselines spanning geolocation, IP, and user-agent signals.

New detector: Suspicious User-Agent Observed in Azure Non-Interactive/SPN Sign In Activity flags authentication activity from offensive tooling (Kali Linux, Parrot Linux, Scout Suite, TruffleHog, and others) in SPN and non-interactive sign-in flows.

New enrichment: Windows Recent Password Change queries Windows Security Event Logs (Event ID 4723 and 4724) to surface last password change time, initiating user, and source computer as lead context.

New scoring layer: Windows Recent Password Change Model applies a configurable confidence modifier to leads based on recency of password reset, reducing noise from post-reset behavior and boosting confidence in potential account-takeover scenarios.

New AXON Threat Hunt Report: Multi-Layer Encoded Command Execution — hunts for nested Base64 encoding/decoding chains in process creation events across Windows, Linux, and macOS as indicators of payload obfuscation.

+4 more

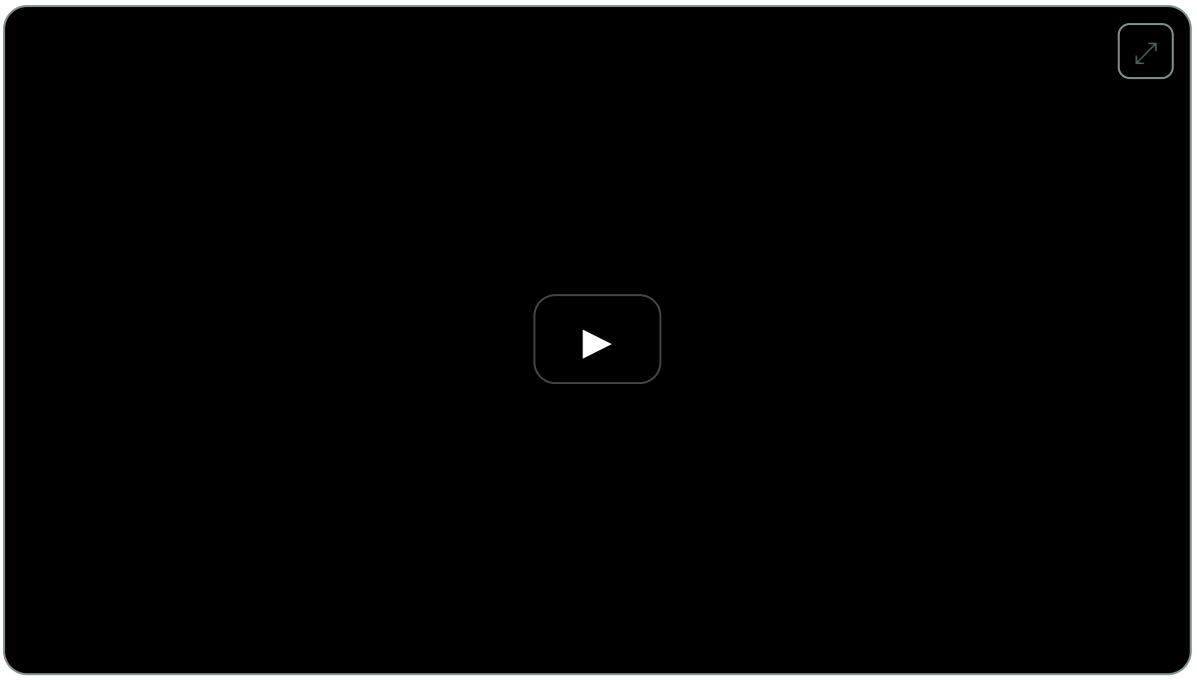
WAS THIS USEFUL? 👍 👎

Microsoft Sentinel Feedback ⓘ

What's new in Microsoft Sentinel NOTES ↗

Microsoft Sentinel is a cloud-native SIEM and SOAR solution that collects, detects, investigates, and responds to security threats across an enterprise.

Sentinel adds data federation, AI playbook GA, custom detection-as-code, UEBA GCP anomalies, and graph-based MCP exploration.



└──▷ **TRY IT**

Deploy custom detection rules directly from your CI/CD pipeline as code, keeping rules version-controlled alongside other security artifacts.

```
$ # After authoring your rule in BICEP, deploy it directly via CLI:
az deployment group create \
  --resource-group <your-rg> \
  --template-file customDetectionRule.bicep \
  --parameters workspaceName=<sentinel-workspace>
```

Access the new consolidated UEBA settings (including Okta V2 and GCP anomaly configuration) from the updated settings page.

📍 In the Microsoft Sentinel portal, go to Settings › Microsoft Sentinel settings › UEBA tab to configure data sources, enable GCP Audit Logs anomaly detections, and manage Okta V2 connector settings.

Adds custom detection rules as code support in Sentinel Repositories (Preview): manage and sync custom detection rules via GitHub or Azure DevOps using the Microsoft Security BICEP extension or BICEP CLI.

Enables linking behavior-based query results from the BehaviorInfo table to new or existing incidents in advanced hunting (Preview), enriching incident context with UEBA data and auto-populating alert metadata and entities.

Introduces a graph tool collection in the Microsoft Sentinel MCP server (Preview) for visual, graph-based exploration of relationships across identities, devices, threats, and signals to assess coverage, dependencies, and configuration gaps.

AI-powered SOAR playbook generator is now generally available (GA), creating Python-based automation workflows through a conversational experience with Cline, an AI coding agent.

Adds a consolidated UEBA settings experience with a new UEBA tab in the Microsoft Sentinel settings page.

+5 more

└─▷ **BREAKING ON UPGRADE**

! Automation rules and Logic Apps playbooks that compare or filter on the `AccountName` property (or `AccountName` in Logic Apps) using the full UPN (e.g., `user@domain.com`) will break after July 1, 2026: `Account Name` will always be set to the UPN prefix only (e.g., `user`) when a full UPN is mapped to Account Name in analytics rule alerts, whereas previously it could be either the full UPN or the prefix.

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Theme and Resolution Options for Dashboard PDF and PNG Exports (Dashboards) (2026-07-10)

12 RELEASES · 2026-06-15 → 2026-07-10 BLOG ↗ LATEST STABLE

Sumo Logic is a cloud-native platform that collects, analyzes, and visualizes machine data and logs for security monitoring, observability, and operational intelligence.

Sumo Logic dashboard PDF/PNG exports gain theme selection, up-to-6000px resolution, and per-panel time range support.

└──▷ HOW TO FIND IT

Generate a high-resolution dark-themed dashboard export for a SOC display wall or executive report.

📍 In the dashboard, go to Export › Export as PDF (or PNG), select Dark theme, set resolution to 6000 pixels, then click Export.

Adds Light or Dark theme selection when exporting dashboards to PDF or PNG and in scheduled reports.

Supports configurable export resolution from 1,500 to 6,000 pixels for sharper output on large displays.

Enables per-panel time range in PDF/PNG exports — panels with overridden time ranges now use their own range instead of the dashboard-level range.

▾ 11 more releases in this issue · 2026-06-15 → 2026-07-10

Include Query and Results in Monitor Email Notifications (Manage) (2026-07-09)

BLOG ↗ STABLE

Sumo Logic monitor email notifications can now embed the triggering search query and result set inline.

└──▷ USE IT

Attach the triggering query and matching results to a monitor's email notification so on-call responders have full context without logging into Sumo Logic.

```json
{
  "notifications": [
    {
      "notification": {
        "connectionType": "Email",
        "recipients": ["soc-team@example.com"],
        "subject": "Monitor Alert: {{Name}}",
        "includeQuery": true,
        "includeResultSet": true
      },
      "runForTriggerTypes": ["Critical", "ResolvedCritical"]
    }
  ]
}
```

Adds `includeQuery` and `includeResultSet` flags to monitor email notifications, embedding the search query and results directly in alert emails.

`Choose the Timestamp Format for Alert Notifications (Manage) (2026-07-08)`

BLOG ↗ STABLE

Sumo Logic adds Notification Settings Standards policy to control timestamp format (Legacy or ISO) in alert payloads.

└──▷ HOW TO FIND IT

Switch alert notification timestamps to ISO 8601 format for easier parsing in downstream SIEM or automation pipelines.

> 📍 In the Sumo Logic console, go to Administration › Policies › Notification Settings Standards and select ISO as the date/time format for alert and recovery notification payloads.

New Notification Settings Standards policy lets org admins choose between Legacy and ISO date/time formats for alert and recovery notification payloads.

`Copy a Dashboard Panel (Dashboards) (2026-07-07)` BLOG ↗ STABLE

Sumo Logic Dashboards now lets you copy any panel to another dashboard without rebuilding the query.

└──▷ HOW TO FIND IT

Reuse a tuned threat-hunting panel across multiple team dashboards without manually recreating the underlying query.

📍 In the Dashboards console, open the source panel menu and select Copy to Dashboard, then choose an existing dashboard or create a new one as the destination.

New Copy to Dashboard action enables copying an existing dashboard panel to another dashboard or a new dashboard, preserving the original query.

Go To Search Bar Enhancements (2026-07-07) BLOG ↗ STABLE

Sumo Logic's Go To search bar now surfaces Library content, Monitors, and SLOs alongside navigation pages.

└──▷ HOW TO FIND IT

Jump directly to a monitor or SLO by name without leaving your current view — use the Go To bar as a universal finder across Library items, Monitors, and SLOs.

📍 In the console, press the Go To search bar (top of the UI), type the name of a dashboard, monitor, or SLO, and select the matching result from the dropdown.

Enables Go To search bar to find dashboards, searches, and folders saved in the Library.

Adds Monitor discovery by name directly from the Go To search bar.

Adds SLO lookup from the Go To search bar without navigating to the SLO page.

Introduces no-match behavior: returns 0 results instead of unrelated suggestions when nothing meets the relevance threshold.

Dashboard Sections Scan Estimate (Dashboards) (2026-07-03) BLOG ↗ STABLE

Sumo Logic Dashboards now show scan estimates at both dashboard and section levels for clearer cost/data visibility.

└──▷ HOW TO FIND IT

Quickly assess total scan cost across an entire dashboard before running heavy queries.

📍 In the dashboard view, click the meter icon (top toolbar) to see estimated scan data split by active and collapsed panels.

Isolate scan cost to a specific dashboard section to identify which panel group is driving high data usage.

📍 In the dashboard view, locate the target section header and click its meter icon to view the scan estimate for panels within that section only.

Adds dashboard-level scan estimates: click the meter icon to view estimated scan data broken down by active and collapsed panels.

Adds section-level scan estimates: click the meter icon on any section header to view estimated scan data for panels within that section.

App Upgrade Email Notifications (Apps) (2026-07-02) BLOG ↗ STABLE

Sumo Logic Next-Gen apps now support email notifications for future app upgrades, with self-service subscribe/unsubscribe.

└──▷ **HOW TO FIND IT**

Subscribe to upgrade notifications when installing a Next-Gen app so your team is alerted to new versions automatically.

> 📍 In the Sumo Logic App Catalog, install a Next-Gen app and opt in to email notifications during the installation flow. To unsubscribe later, navigate to the installed app in the App Catalog and select the unsubscribe option.

Adds opt-in email notifications for Next-Gen app upgrades, including app name, new version, release notes, and a direct App Catalog link.

Enables users with the Manage Apps capability to subscribe or unsubscribe from upgrade notifications for apps installed by others.

Automatically removes upgrade notifications when a Next-Gen app is uninstalled.

New user onboarding now defaults to the New UI (2026-07-02) BLOG ↗ STABLE

New Sumo Logic users and orgs now onboard directly into the New UI by default.

└──▷ **HOW TO FIND IT**

New users who prefer the previous layout can revert immediately after signing up.

> 📍 In the New UI, open the left navigation menu and select Switch to classic UI.

New UI is now the default onboarding experience for all new Sumo Logic users and organizations; Classic UI remains accessible via `Switch to classic UI` in the left navigation menu.

Apps, Solutions, and Collection Integrations - June Release (2026-06-30)

BLOG ↗ STABLE

Sumo Logic adds Amazon Overview app, Next-Gen app migrations, and hashing support in 14 updated Source Templates.

New Amazon Overview app delivers unified visibility into AWS service activity, performance, and health metrics across your cloud environment.

Adds new dashboards, monitor alerts, and an IDP category to Active Directory 2012+ (OpenTelemetry), Active Directory 2012+ (JSON), Auth0, Azure Active Directory, JumpCloud Directory Insights, and OneLogin apps.

Adds a `custom_gpt` selection variable to the ChatGPT Compliance app dashboard for easier filtering, alongside updated queries aligned to the latest log format.

Introduces hashing support for fields in filelog receiver-based Source Templates, enabling sensitive field obfuscation at collection time.

Adds field table validation to filelog receiver-based Source Templates to catch misconfigurations before data is collected.

+2 more

└──▷ **BREAKING ON UPGRADE**

! The deprecated batch processor has been removed from the OpenTelemetry configuration for Active Directory 2012+ - OpenTelemetry; the updated configuration requires Sumo Logic OpenTelemetry Agent version 0.153 or later.

! The deprecated batch processor configurations have been removed from 14 Source Templates; setups using those configurations must be updated to ensure compatibility with current collector versions.

Unified Threat Intelligence Configuration and Management (2026-06-26) BLOG ↗

STABLE

Sumo Logic unifies Threat Intelligence source configuration and indicator uploads into a single Data Management hub.

└──▷ **HOW TO FIND IT**

Onboard a new threat intel source and review its indicator count from the unified console — no page-switching required.

📍 In the console, go to Data Management › Threat Intelligence, click Add Source, select your threat intelligence source, choose a collector, and complete the configuration wizard.

Push a custom set of indicators (e.g., IOCs from an incident) directly into Threat Intelligence without a formal source integration.

📍 In the console, go to Data Management › Threat Intelligence, then use the manual upload indicators option to import your indicator file.

Centralizes Threat Intelligence management under Data Management > Threat Intelligence, replacing fragmented multi-page navigation.

New Add Source flow lets analysts select a threat intel source, choose a collector, and complete configuration in one guided workflow.

New manual indicator upload option allows practitioners to add indicators directly without a separate source integration.

Updated Sources listing page surfaces Status, Source Name, Description, Storage Consumed, and Number of Indicators at a glance.

Manage Features for Child Organizations (MSSP) (2026-06-24) BLOG ↗ STABLE

MSSP admins can now opt child orgs in/out of Sumo Logic features directly from the Organizations UI.

└──▷ **HOW TO FIND IT**

Enable or disable a specific Sumo Logic feature for an individual child organization without contacting support.

📍 In the Sumo Logic console, go to Administration › Organizations, select a child organization, then use the Manage Features option to toggle individual features on or off.

Sumo Logic Deployment in AWS European Sovereign Cloud Region (2026-06-15)

BLOG ↗ STABLE

Sumo Logic launches a new service region in AWS European Sovereign Cloud for strict EU data residency and compliance.

New Sumo Logic deployment region in AWS European Sovereign Cloud supports EU data residency, sovereignty controls, and compliance requirements.

Enables organizations operating under strict European regulatory frameworks to run log analytics, security monitoring, and observability workloads with reduced latency.

WAS THIS USEFUL? 👍 👎

Google SecOps

Feedback ⓘ July 01, 2026 (2026-07-01)

6 RELEASES · 2026-06-16 → 2026-07-01 BLOG ↗ LATEST STABLE

Google SecOps is a cloud-based security operations platform that helps organizations detect, investigate, and respond to cybersecurity threats using advanced analytics and intelligence.

Google SecOps introduces Security Tokens for metering agentic SOC consumption across GA agents.

Adds Security Tokens as a new metering unit for tracking and billing consumption by generally available security agents.

Supports agent invocation via web interface, CLI, chat, and Model Context Protocol (MCP) — all metered through Security Tokens.

Excludes assistive features (standard chat panels, automated summaries) and preview agents from Security Token consumption.

▾ 5 more releases in this issue · 2026-06-16 → 2026-07-01

June 30, 2026 (2026-06-30) BLOG ↗ STABLE

Google SecOps raises event rule limits and launches unified rules interface in public preview.

└──▷ HOW TO FIND IT

Access the new unified rules interface to manage both custom and curated detection rules from a single dashboard.

📍 In the Google SecOps console, navigate to Detection › Rules. The unified rules interface loads by default in public preview. To revert, click Switch to the legacy experience at the top right of the screen.

Increases multiple event rule limits to 200 for Enterprise customers and 400 for Enterprise+ customers.

Introduces unified rules interface (public preview) combining custom and curated rule management into a single workflow with a redesigned dashboard, advanced rule editor, and expanded API capabilities.

June 28, 2026 (2026-06-28) BLOG ↗ STABLE

Google SecOps adds 15 new log types including MCP, OT, and AI platform sources, plus updates 130+ default parsers.

└──▷ HOW TO FIND IT

Enable ingestion of Model Context Protocol logs to detect and investigate AI agent activity (dev, modify, and view actions) in your environment.

> 📍 In the Google SecOps console, go to Settings › Log Ingestion, search for log types MCPDEV, MCPMODIFY, or MCPVIEW, and configure the appropriate feed for your MCP deployment.

Onboard NetApp Ransomware Resilience telemetry to correlate storage-layer ransomware signals with endpoint and network detections.

> 📍 In the Google SecOps console, go to Settings › Log Ingestion, search for `NETAPP_RANSOMWARE_RESILIENCE`, and configure a feed pointing to your NetApp data source.

Ingest Reco AI CSPM findings to surface cloud security posture issues alongside other SecOps detections.

> 📍 In the Google SecOps console, go to Settings › Log Ingestion, search for `RECO_CSPM`, and configure the feed with your Reco AI credentials and endpoint.

Adds new log type `CISCO_SECURE_ACCESS_ENROLLMENT` for Cisco Secure Access Enrollment ingestion.

Adds new log type `CISCO_SECURE_ACCESS_NETWORK` for Cisco Secure Access Network ingestion.

Adds new log type `CYBERARK_CERTIFICATE_MANAGER_SAAS` for CyberArk Certificate Manager SaaS audit data.

Adds new log type `GEMINI_ENTERPRISE_AGENT_PLATFORM` for Gemini Enterprise Agent Platform activity.

Adds new log type `GEOSCADA_OT` for Schneider Electric GeoScada operational technology logs.

+9 more

June 23, 2026 (2026-06-23) BLOG ↗ STABLE

Google SecOps adds Gemini Cloud Assist directly in Feed Management for guided feed setup and troubleshooting.

└──▷ **HOW TO FIND IT**

Get AI-guided help configuring a new data feed without leaving the Feed Management UI.

> 📍 In Google SecOps, go to Feed Management, then click the Ask Gemini Cloud Assist button to open the Gemini Cloud Assist panel and ask questions about feed configuration, ingestion prerequisites, or setup issues.

June 17, 2026 (2026-06-17) BLOG ↗ STABLE

Google SecOps adds configurable auto-collapse for the query editor to maximize search results screen space.

└──▷ HOW TO FIND IT

Maximize visible search results during investigations by collapsing the query editor automatically after each query runs.

📍 In the console, go to the query editor settings and enable the Auto-collapse option so the editor collapses after you run a search.

New Auto-collapse setting for the query editor automatically hides the editor after a search runs, giving more screen real estate to results.

June 16, 2026 (2026-06-16) BLOG ↗ STABLE

Google SecOps launches a monthly documentation changelog to track major doc updates.

└──▷ HOW TO FIND IT

Stay current on major Google SecOps documentation changes by bookmarking the new monthly changelog.

📍 In the Google SecOps docs, navigate to Release Notes › Documentation changelog to view the monthly summary of major documentation updates.

WAS THIS USEFUL? 👍 👎

Exabeam

Feedback ⓘ New Threat Center Endpoints (2026-06-16) BLOG ↗

Exabeam is a security information and event management (SIEM) platform that uses AI and behavioral analytics to detect, investigate, and respond to cybersecurity threats.

Exabeam adds Threat Center API endpoints to retrieve threat timelines by alert ID or case ID.

New API endpoint to retrieve a threat timeline given a specific alert ID.

New API endpoint to retrieve a threat timeline given a specific case ID.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Sophos Central

Feedback ⓘ snapshot-20260713 NOTES ↗

Sophos Central is a cloud-based security management platform that unifies endpoint, network, and cloud protection with centralized threat detection and response.

Sophos Central adds DNS Protection APIs, Gen-AI web filtering, MFA enforcement, new email tiers, and expanded MDR/XDR capabilities.

└──▷ TRY IT

Rotate API credentials before they expire — use the new expiration alerts as your trigger to call the DNS Protection API and refresh keys.

```
$ curl -X GET 'https://api.central.sophos.com/dns-
protection/v1/policies' \
  -H 'Authorization: Bearer <access_token>' \
  -H 'X-Tenant-ID: <tenant_id>'
```

Enable Newly Registered Domain protection in Sophos Email to block phishing attempts from freshly created domains before they reach users.

📍 In Sophos Central, go to My Products › Sophos Email › Policies, open or create an Inbound policy, and enable the `Newly Registered Domain (NRD) Protection` option under threat detection settings.

Apply a Generative AI category block in Protected Browser to prevent data exfiltration via consumer AI tools.

📍 In Sophos Central, go to My Products › Workspace Protection › Protected Browser › Policies, open your policy, navigate to App Controls, use the new categorized app picker to select the `Generative AI` category, and set the action to Block.

Adds Web Control filtering for Generative AI sites with more granular per-site control and simplified policy management.

Enables alerts for Account Health Check (AHC) issues, gradually rolling out to surface unaddressed configuration risks.

Adds categorized Gen-AI app picker in Protected Browser policy creation, plus reporting on file uploads and downloads for risky data flow visibility.

Consolidates DNS Protection Endpoint and DNS filtering policies onto a single management page for Workspace Protection customers.

Requires administrators to register at least two MFA methods for sign-in to reduce lockout risk.

+24 more

WAS THIS USEFUL?  

Huntress

Feedback ⓘ

AI Search for Huntress Managed SIEM has officially gone GA

9 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Huntress is a managed security platform that detects, investigates, and remediates cyber threats for small and mid-sized businesses through persistent endpoint monitoring.

Huntress Managed SIEM gains GA AI Search, enabling plain-English log queries with automatic ESQL translation.

└──▷ HOW TO FIND IT

Run a plain-English investigation query against SIEM logs, then inspect the generated ESQL to learn query syntax or refine further.

📍 In Huntress Managed SIEM, enter your query in the AI Search bar (e.g. "show me failed logins from external IPs in the last 24 hours"), press Search, then click the ESQL button to view and copy the translated ESQL query.

Adds AI Search (GA) to Huntress Managed SIEM, allowing natural-language queries against SIEM logs without requiring ESQL or the Query Builder.

Enables on-demand ESQL translation: after a plain-English search, clicking the ESQL button reveals the generated ESQL query to accelerate analyst upskilling.

▾ 8 more releases in this issue · seen 2026-07-13

Additional Security Controls in Managed ISPM NOTES ↗ STABLE

Huntress Managed ISPM adds SharePoint/OneDrive/Teams security controls and expands Continuous Enforcement to five new policies.

└──▷ HOW TO FIND IT

Enable Continuous Enforcement on the 'Require MFA to register or join devices' policy to automatically detect and remediate drift without manual intervention.

📍 In the Huntress portal, go to Managed ISPM › Security Controls, locate 'Require MFA to register or join devices', and toggle on Continuous Enforcement to enable drift detection and auto-remediation.

New SharePoint Online control: enforces idle session timeout for SharePoint and OneDrive.

New SharePoint Online control: blocks anonymous sharing links in SharePoint and OneDrive.

New SharePoint Online control: prevents users from self-creating SharePoint sites.

New SharePoint Online control: retains deleted user OneDrive content for at least 90 days.

New Microsoft Teams control: blocks communication with unmanaged Teams tenants.

+2 more

Huntress Managed EDR for macOS now ejects deceptive installers NOTES ↗

STABLE

Huntress Managed EDR for macOS auto-ejects malicious DMGs before payload execution, blocking infostealers.

Automatically ejects deceptive/malicious DMG installers on macOS before their payloads can execute, blocking infostealer malware.

Delivers on-screen user notifications explaining why a DMG was ejected.

Triggers a Huntress SOC investigation on every DMG ejection event to confirm endpoint cleanliness.

Huntress Managed EDR stops macOS infostealers before they run NOTES ↗

STABLE

Huntress Managed EDR now blocks macOS infostealers at the Gatekeeper-override moment on macOS 14+.

New macOS infostealer prevention: detects and stops infostealer execution at the moment a user is socially engineered into overriding Apple's Gatekeeper, available on macOS 14 and later.

Lock down Microsoft 365 identities today as Managed ISPM hits General Availability

NOTES ↗ STABLE

Huntress Managed ISPM hits GA, delivering continuous M365 identity hardening across Entra ID, Exchange, SharePoint, and Teams.

└──▷ HOW TO FIND IT

Use Learning Mode to model the blast radius of a new Conditional Access policy before it goes live — preventing unexpected user lockouts across tenants.

📍 In the Huntress console, navigate to ISPM › Conditional Access Policies, select a recommended template, then enable Learning Mode to preview impact before deploying enforcement.

Adds Managed ISPM (Identity Security Posture Management) at General Availability, continuously hardening Microsoft 365 environments against misconfigurations and over-permissioned accounts.

New Huntress-managed security controls span Microsoft Entra ID, Exchange Online, SharePoint, and Teams — covering the most commonly abused identity attack surfaces.

Adds Conditional Access policy management with recommended templates, giving practitioners a curated baseline for CA enforcement.

New Learning Mode delivers pre-deployment impact analysis for Huntress-managed Conditional Access policies before any enforcement is applied.

Managed Deployment enables a continuously updated identity framework that automatically keeps all managed organizations aligned with the latest security updates.

`+1 more`

`OSINT 2 Simulation is in GA` **NOTES ↗** STABLE

Huntress adds OSINT 2 Simulation in GA, teaching learners the risks of oversharing online.

New OSINT 2 Simulation (fifth simulation) now generally available, covering risks of oversharing personal information online.

Accompanies July's managed learning assignment for structured security awareness training.

`SAT Custom HTML Phishing is GA` **NOTES ↗** STABLE

Huntress Managed SAT now offers GA Custom HTML Phishing Scenarios for targeted spearphishing simulations.

`SIEM Okta integration is now GA` **NOTES ↗** STABLE

Huntress Managed SIEM adds Okta as a GA log source with SOC-built detections for identity-based attacks.

`└──▷` **HOW TO FIND IT**

Connect Okta System Log to Huntress Managed SIEM to enable identity-layer detection coverage.

📍 In the Huntress portal, navigate to SIEM › Log Sources › Add Source, select Okta (API – Okta System Log), and follow the configuration guide at `https://support.huntress.io/hc/en-us/articles/50820786591379-API-Okta-System-Log`

Adds Okta as a supported identity and authentication log source in Huntress Managed SIEM (generally available).

New SOC-built detections for Okta covering credential attacks, privilege escalation, MFA bypass and fatigue, account takeover, and federated identity manipulation.

Extends SIEM visibility to the identity layer for organizations using Okta as their IdP or SSO provider.

The Microsoft Defender for Endpoint (MDE) integration now supports GCC High tenants

NOTES ↗ STABLE

Huntress MDE integration now supports GCC High tenants, routing API traffic to .us endpoints automatically.

Supports onboarding Microsoft Defender for Endpoint (MDE) tenants in GCC High environments, with login, Graph, and MDE API traffic automatically routed to GCC High .us endpoints.

Enables GCC High tenants already mapped in ITDR to be onboarded into MDE using the same tenant.

WAS THIS USEFUL?  

Trend Vision One

Feedback ⓘ

Zero Trust Secure Access On-premises Gateway version 1.0.87 (2026-07-13)

14 RELEASES · 2026-06-26 → 2026-07-13 BLOG ↗ LATEST STABLE

Trend Vision One is a unified cybersecurity platform that provides extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

On-premises Gateway 1.0.87 adds X-Authenticated-User header support for upstream/downstream proxy auth in Forward proxy mode.

└──▷ HOW TO FIND IT

When a downstream proxy already authenticates users and stamps X-Authenticated-User, enable the gateway to accept that header as identity — eliminating the need for a second credential challenge in Forward proxy mode.

> 📍 In the Vision One console, go to Zero Trust Secure Access › On-premises Gateway › Forward Proxy Settings and enable `Accept X-Authenticated-User header from downstream proxy` as the authentication source.

When the gateway needs to pass authenticated user identity to an upstream proxy, configure it to stamp the X-Authenticated-User header on outbound requests so the upstream proxy can enforce per-user policy.

> 📍 In the Vision One console, go to Zero Trust Secure Access › On-premises Gateway › Forward Proxy Settings, configure the upstream proxy, and enable `Stamp X-Authenticated-User header to upstream proxy`.

Supports parsing the X-Authenticated-User header stamped by a downstream proxy as an alternative authentication method in Forward proxy mode.

Supports stamping the X-Authenticated-User header onto configured upstream proxy requests in Forward proxy mode.

Updates DLP SDK and built-in compliance templates with new content.

Adds protocol compatibility with web servers that return HTTP 103 Early Hints responses.

▾ 13 more releases in this issue · 2026-06-26 → 2026-07-13

> TrendAI Vision One Endpoint Security agent now supports Red Hat Enterprise Linux 8 Pow
>
> BLOG ↗ STABLE
>
> Trend Vision One Endpoint Security agent now supports Red Hat Enterprise Linux 8 PowerPC (ppc64le) deployments.

Supports deploying the Endpoint Security agent on Red Hat Enterprise Linux 8 PowerPC (ppc64le) endpoints.

Enables Server & Workload Protection coverage on RHEL 8 ppc64le systems.

Enables XDR for Endpoint (EDR) telemetry and detection on RHEL 8 ppc64le systems.

`Custom attack prompts for AI Scanner (2026-07-07)` **BLOG ↗** **STABLE**

Trend Vision One AI Scanner adds custom attack prompt support for deterministic, repeatable AI model security testing.

└──▷ HOW TO FIND IT

Reach the custom prompts feature in the console to supply your own jailbreak payloads and tag them for OWASP/CVSS compliance reporting.

> 📍 In the console, go to AI Security → AI Application Security → Custom prompts, then add your hand-crafted attack payloads, define evaluation criteria per objective, and apply CVSS/OWASP framework tags before running the scan.

Supports custom hand-crafted attack payloads in AI Scanner alongside the built-in attack library, enabling testing of known jailbreaks, regression cases, and novel attack variations.

Runs custom prompts exactly as written with no automatic technique or modifier applied, ensuring deterministic and repeatable results.

Enables user-defined evaluation criteria per objective for custom prompt tests.

Adds CVSS and OWASP framework tagging for custom prompts so results appear in compliance reporting alongside built-in test results.

`Agentless Vulnerability & Threat Detection version 1.198.07031506 (2026-07-06)`

BLOG ↗ **STABLE**

Agentless Vulnerability & Threat Detection adds configurable scan scheduling for Microsoft Azure environments.

└──▷ HOW TO FIND IT

Reach the Azure scan scheduling controls to set frequency and start time, or toggle scanning off without a Terraform redeployment.

> 📍 In the console, go to Cloud Security → Cloud Accounts → Azure, then configure scan frequency, start time, or enable/disable scanning for your subscription.

Enables configuration of scan frequency and start time for Agentless Vulnerability & Threat Detection in Azure subscriptions via Terraform.

Supports enabling or disabling Azure scanning on-demand without redeploying the Terraform template.

TrendAI Vision One has a new Indonesia Data Center (2026-07-01) BLOG ↗

STABLE

Trend Vision One opens an Indonesia data center, giving regional users a local service endpoint.

└──▷ **HOW TO FIND IT**

Update firewall allowlists to permit traffic to the new Indonesia-region FQDNs after migrating your tenant to the local data center.

📍 In the console, go to the Indonesia - firewall exceptions page and update your firewall rules to allow the listed Indonesia-region service FQDNs.

Back up and restore Deep Discovery Inspector appliance configurations (2026-06-29)

BLOG ↗ STABLE

Trend Vision One adds centralized backup and restore for Deep Discovery Inspector appliance configs via the console.

└──▷ **HOW TO FIND IT**

Restore a DDI appliance to a known-good configuration after a misconfiguration — without SSHing into the appliance directly.

📍 In the console, go to Network Security → Network Inventory, select the target Deep Discovery Inspector appliance, open System Maintenance, choose a backup from the backup history, and select Restore.

New System Maintenance screen in the Trend Vision One console enables creating, downloading, deleting, and restoring configuration backups for connected Deep Discovery Inspector appliances without logging into each appliance individually.

Supports backup history management — view past backups and restore an appliance from any previously saved configuration snapshot.

Requires Deep Discovery Inspector 6.9 or later; accessible via Network Security → Network Inventory.

Data Loss Prevention and AI content detection logging (2026-06-29) BLOG ↗

STABLE

Zero Trust Secure Access now logs DLP and AI content snippets (up to 100 bytes) including prompt injection detections.

└──▷ **HOW TO FIND IT**

Enable DLP and AI content snippet logging to capture evidence of policy-matching traffic and prompt injection attempts in Zero Trust Secure Access.

📍 In the console, go to Zero Trust Secure Access → Internet Access and AI Secure Access Configuration → Global Settings → DLP/AI detection logs, then enable content snippet logging for DLP profiles, AI content inspection profiles, and prompt injection detections.

Enables content snippet logging (up to 100 bytes) for traffic matching a Data Loss Prevention profile via Internet Access and AI Secure Access.

Enables content snippet logging for traffic matching an AI content inspection profile.

Adds prompt injection detection logging within Zero Trust Secure Access Internet Access and AI Secure Access flows.

Email notifications available for detection model updates (2026-06-29) BLOG ↗

STABLE

Trend Vision One adds email notifications for Detection Model Management changes — new, updated, or deprecated models.

└──▷ HOW TO FIND IT

Subscribe to detection model change alerts so your team is immediately notified when a model is deprecated or a new detection is available.

📍 In the console, go to Administration → Notifications, then click Detection model updates to configure email notification subscriptions.

File Security Storage template version 4.2.0 for AWS (2026-06-29) BLOG ↗

STABLE

Trend Vision One File Security Storage AWS template v4.2.0 adds support for S3 KMS Encryption.

Supports S3 KMS Encryption for protected S3 buckets scanned by File Security Storage.

Internet Access rule simulator (2026-06-29) BLOG ↗ STABLE

Trend Vision One adds an Internet Access rule simulator to test and validate Zero Trust Secure Access rules before deployment.

└──▷ HOW TO FIND IT

Validate a new Internet Access rule to catch gaps or conflicts before it affects live traffic.

📍 In the console, go to Zero Trust Secure Access → Secure Access Rules → Internet Access, then open the rule simulator to test your rule configuration before deployment.

`Threat actor attribution in sandbox reports (2026-06-29)` BLOG ↗ STABLE

Sandbox Analysis reports now surface threat actor and campaign attribution via a new Associated threats tab.

└──▷ **HOW TO FIND IT**

After submitting a suspicious file to Sandbox Analysis, check attribution without pivoting to a separate tool.

📍 In Threat Intelligence → Sandbox Analysis, open a completed object profile and select the Associated threats tab to view matched threat actor names and linked campaign reports.

Adds an Associated threats tab to sandbox object profiles, correlating submitted files/IOCs with known threat actors and emerging threats from Threat Intelligence Hub.

Enables direct links from sandbox reports to full threat actor and campaign reports in Threat Intelligence Hub, eliminating separate intelligence lookups.

`Zero Trust Secure Access On-premises Gateway version 1.0.86 (2026-06-29)`

BLOG ↗ STABLE

ZTSA On-premises Gateway 1.0.86 adds OCR-based DLP, Grok inspection, and M365 Copilot Tenant Restriction V2 support.

Adds OCR integration to DLP and AI content inspection, enabling text extraction and policy enforcement on image-based content.

Supports Microsoft 365 Tenant Restriction V2 for controlling Copilot access across tenants.

Adds content inspection support for Grok services.

`Zero Trust Secure Access adds support for tenancy restrictions for Copilot (2026-06-29`

BLOG ↗ STABLE

Zero Trust Secure Access now enforces tenancy restrictions for Microsoft Copilot, blocking personal-account logins from corporate networks.

└──▷ **HOW TO FIND IT**

Prevent employees from accessing Microsoft Copilot with personal accounts (Gmail, Outlook) over the corporate network by enabling a tenancy restriction rule.

📍 In the console, go to Zero Trust Secure Access → Secure Access Resources → Tenancy Restrictions and create or edit a rule targeting Microsoft Copilot, setting the allowed tenant to your corporate domain.

Adds tenancy restriction support for Microsoft Copilot in Internet Access and AI Secure Access rules, limiting access to corporate-domain accounts only and blocking personal accounts (e.g., Gmail, Outlook) from the corporate network.

Security Playbooks collects forensic evidence (2026-06-26) BLOG ↗ STABLE

Security Playbooks gains a native Collect Evidence action node to trigger forensic collection on endpoints without playbook chaining.

└──▷ **HOW TO FIND IT**

Automatically collect forensic evidence on high-severity incidents without manual steps or chaining — results appear in the Forensics workspace alongside the Workbench timeline.

📍 In the console, go to Workflow and Automation → Security Playbooks, create or edit a Workbench insight progression update playbook, add a condition node (e.g., Insight score › 80 or Severity = High), then add a Collect Evidence action node as the next step and save.

Adds a Collect Evidence action node to Security Playbooks, enabling automated forensic evidence collection on impacted Windows and Linux endpoints directly within a playbook.

Supports triggering evidence collection from endpoint response, automated response, and Workbench insight progression update playbook types.

Enables criteria-based automation (e.g., severity is high or Insight score exceeds 80) to kick off forensics without manual intervention or chaining separate playbooks.

Surfaces collected evidence alongside Workbench events and in the Forensics workspace, consolidating incident timelines into a single workflow.

WAS THIS USEFUL? 👍 👎

Automox

Feedback ⓘ

Introducing Remote Tools: Remote Command & File Transfer (2026-06-29)

3 RELEASES · 2026-06-16 → 2026-06-29 BLOG ↗ LATEST STABLE

Automox is a cloud-native endpoint management platform that automates patching, configuration, and software deployment across Windows, macOS, and Linux devices.

Automox adds Remote Command and File Transfer to its new Remote Tools umbrella, enabling shell access and file I/O without a full desktop session.

.
.
.

└──▷ HOW TO FIND IT

Run a diagnostic command on a managed endpoint during an incident without interrupting the end user's active session.

> 📍 In the Automox console, navigate to a managed device › Remote Tools › Remote Command, open an interactive shell session, and run your diagnostic commands in real time.

Push a remediation script or pull back logs from a device without starting a full screen-sharing session.

> 📍 In the Automox console, navigate to a managed device › Remote Tools › File Transfer, then upload a script or download log files as a targeted action.

Introduces Remote Tools — a unified remote support hub that consolidates Remote Control, Remote Command, and File Transfer in one place.

New Remote Command capability opens an interactive shell session on a managed Windows or macOS device in the background, without taking over the end user's screen.

New File Transfer capability lets practitioners push scripts, configs, or installers to — and pull logs or diagnostic output from — a managed device without launching a full remote control session.

▾ 2 more releases in this issue · 2026-06-16 → 2026-06-29

> Slack for Windows has moved to MSIX - how Automox supports the transition (2026-06-19)
>
> BLOG ↗ STABLE
>
> Automox adds native Slack MSIX detection and patching plus a migration Worklet for Windows MSI-to-MSIX transitions.

└──▷ **HOW TO FIND IT**

Migrate Windows endpoints from legacy Slack MSI or EXE installations to MSIX in a single controlled action without waiting for a scheduled patch window.

> 📍 In the Automox console, go to Worklets, search for `Windows - Software Lifecycle - Migrate Slack to MSIX`, assign it to the target device group, then click Run Now.

Adds native detection and patching support for Slack MSIX packages on Windows, reported under the existing Slack catalog title with no policy reconfiguration required.

Supports Slack MSIX on both x64 and ARM64 Windows devices.

Introduces the base Worklet `Windows - Software Lifecycle - Migrate Slack to MSIX`, available to all customers (including non-Worklet Catalog subscribers), to remove legacy MSI/EXE installs and provision the current MSIX package in a single controlled action.

> └──▷ **BREAKING ON UPGRADE**
>
> ! Existing Slack MSI installs on Windows will no longer receive newer MSI updates through routine Automox patching, as Automox has deprecated the legacy MSI as a patch target following Slack's retirement of its Windows MSI installers.

`Target patching by real-world risk: KEV & EPSS now on Windows OS and in Advanced Patch`

BLOG ↗ STABLE

Automox adds KEV & EPSS exploit-risk targeting to Windows OS patching and Advanced Patch policies.

.
.

└──▷ **HOW TO FIND IT**

Automatically patch Windows OS CVEs that are actively exploited in the wild without manual triage, by targeting KEV-flagged vulnerabilities in an Advanced Patch policy.

> 📍 In the Automox console, go to Policies › Advanced Patch › create or edit a policy, then under Patch Filters enable `KEV` to scope patching to CISA Known Exploited Vulnerabilities on Windows OS endpoints.

Prioritize Windows OS patches most likely to be exploited in the next 30 days by filtering on EPSS score within an Advanced Patch policy.

> 📍 In the Automox console, go to Policies › Advanced Patch › create or edit a policy, then under Patch Filters enable `EPSS` and set a probability threshold to automatically target high-risk Windows OS packages.

> Extends KEV (CISA Known Exploited Vulnerabilities) and EPSS scoring to Windows OS packages, visible on Software and Device Details pages.
>
> Enables automated remediation targeting by KEV and EPSS status inside Advanced Patch policies, eliminating manual CVE triage.
>
> Surfaces exploited and high-probability-of-exploitation vulnerabilities alongside third-party titles for unified risk-based prioritization.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Node v6.6.0: Expanded ClickHouse type support (2026-07-10)

12 RELEASES · 2026-06-15 → 2026-07-10 BLOG ↗ LATEST STABLE

Tenzir Changelog is an RSS feed tracking updates, new features, and fixes for Tenzir's open-source security data pipeline and threat detection platform.

Tenzir v6.6.0 expands ClickHouse integration with LowCardinality/DateTime64 support and adds a reference date for partial parse_time formats.

Adds reference date support for partial `parse_time` format strings, enabling relative time parsing without a full timestamp.

Supports additional ClickHouse column types in `to_clickhouse`, including LowCardinality and DateTime64.

Enables `to_clickhouse` to tolerate ClickHouse columns that have default values, avoiding pipeline failures on schema mismatches.

Adds warnings in `to_clickhouse` instead of silently dropping events when column mapping issues occur.

▾ 11 more releases in this issue · 2026-06-15 → 2026-07-10

> Tenzir Platform v1.35.0: Multi-Factor Authentication (2026-07-09) BLOG ↗ STABLE
>
> Tenzir Platform v1.35.0 adds optional MFA support and improves pipeline Diagnostics tab performance.
>
> Adds optional multi-factor authentication (MFA) support for the public Tenzir Platform.
>
> Improves performance of the pipeline Diagnostics pane.
>
> Enables graceful shutdown for Docker Compose nodes.

> Tenzir Platform v1.36.0: Inline UDO Expansion in the Explorer (2026-07-09) BLOG ↗ STABLE
>
> Tenzir Platform v1.36.0 adds inline UDO expansion in the Explorer, revealing operator bodies with full TQL syntax highlighting.
>
> └──▷ HOW TO FIND IT
>
> When debugging a pipeline that calls a UDO, expand it inline in the Explorer to see the operator's full TQL body without leaving the editor.

📍 In the Explorer, open a pipeline containing a UDO call, then click the inline expand control next to the UDO call to reveal its body with TQL syntax highlighting, including any nested UDO calls.

Adds inline expansion of User-Defined Operators (UDOs) in the Explorer, letting practitioners expand any UDO call in place to inspect its full TQL body, including nested UDO calls.

Tenzir Ship v1.11.0: GitHub release title formatting (2026-07-07)

BLOG ↗

STABLE

Tenzir Ship v1.11.0 adds project name and version to GitHub release titles with custom format support.

GitHub release titles now include the project name, version, and release title by default, keeping untitled releases clean.

Supports custom release title formats via the CLI and Python API.

Tenzir Skills v3.0.0: Streamlined Tenzir skill collection (2026-07-05)

BLOG ↗

STABLE

Tenzir Skills v3.0.0 renames the docs skill to `tenzir` and removes the shared review workflow.

Renames the bundled documentation skill from `tenzir-docs` to `tenzir` (breaking change — update any references).

Removes the opinionated shared review workflow from the bundled skill collection.

Stabilizes scheduled ASIM updates so that upstream branch movement without content changes no longer triggers unnecessary syncs.

Tenzir Node v6.5.0: Reliability fixes and faster summarize (2026-07-03)

BLOG ↗ STABLE

Tenzir Node v6.5.0 adds map_keys(), expanded metrics, assert_throughput bounds, and hive-partitioned sink column pruning.

Adds `map_keys` function to extract keys from record-type fields in pipeline expressions.

Adds ingress/egress metrics to additional operators, broadening pipeline observability.

Supports upper bounds in `assert_throughput` for capping expected throughput ranges.

Hive-partitioned filesystem sinks now omit partition columns from written files, avoiding data duplication.

Faster `summarize` throughput on interleaved-key inputs.

Tenzir Ship v1.10.0: Improved PR metadata handling (2026-07-03) BLOG ↗

STABLE

Tenzir Ship v1.10.0 adds PR metadata validation controls for changelog projects.

└──▷ **USE IT**

Enforce that every unreleased entry carries a PR reference, catching omissions before merge.

```yaml
require_pr: true
```

Prevent PR and author fields from appearing in entries when your project policy strips them.

```yaml
omit_pr: true
omit_author: true
```

New `require_pr` option warns when an unreleased changelog entry is missing its pull request reference.

Validation now rejects PR and author metadata on entries when the project config sets `omit_pr` or `omit_author`, preventing stray metadata from slipping in.

Agent skill updated to respect `omit_pr` and `omit_author` project configuration options.

Tenzir Node v6.4.0: Package constants and ClickHouse JSON (2026-06-29) BLOG ↗

STABLE

Tenzir v6.4.0 adds package constants, ClickHouse JSON column support, and embedded platform control endpoints in tokens.

Adds package-level constants that any pipeline within the package can reference, enabling reusable configuration across pipelines.

Supports native ClickHouse JSON columns in the `to_clickhouse` operator for richer schema flexibility when writing to ClickHouse.

Embeds platform control endpoints directly in node tokens, simplifying node onboarding without separate endpoint configuration.

Tenzir Node v6.3.0: Graceful pipeline shutdown (2026-06-24)

BLOG ↗ STABLE

Tenzir v6.3.0 adds graceful pipeline shutdown, ClickHouse source, HTTP proxy support, and automatic input format detection.

Adds graceful pipeline shutdown with data draining — in-flight events are flushed before a pipeline or node stops, eliminating event loss on shutdown.

New `from_clickhouse` source operator enables reading data directly from ClickHouse databases.

Supports configuring an HTTP proxy for outbound operators.

Automatic input format detection removes the need to specify formats manually for incoming data.

Adds call-site annotations in diagnostics from user-defined operators, tracing errors back to the top-level pipeline for easier debugging.

+4 more

Tenzir Platform v1.34.0: OpenTelemetry tracing and Stream columns (2026-06-22)

BLOG ↗ STABLE

Tenzir Platform v1.34.0 adds OpenTelemetry frontend tracing and dedicated Stream view columns for common fields.

Adds dedicated columns for common fields (e.g., timestamps) in the Stream view for faster at-a-glance triage.

Adds OpenTelemetry tracing for the Platform frontend, enabling end-to-end visibility into backend request handling.

Removes direct WebSocket Gateway connection support — operators relying on it must migrate to the new connection path.

Tenzir Helm Charts v0.1.0: Initial Tenzir Node Helm chart (2026-06-18)

BLOG ↗ STABLE

Tenzir Node Helm chart debuts with hardened defaults, flexible networking, and checksum-driven rollouts.

New Helm chart for deploying Tenzir Node on Kubernetes.

Supports per-node configuration with checksum-driven rollouts to trigger pod restarts on config changes.

Two patterns for exposing listener ports to suit different networking topologies.

Hardened container security defaults applied out of the box.

Optional NetworkPolicy and PodDisruptionBudget resources for production-grade

deployments.

Tenzir Skills v2.2.0: Security schema reference skills (2026-06-15) BLOG ↗

STABLE

Tenzir Skills v2.2.0 adds reference skills for 7 major security schemas: ASIM, Splunk CIM, ECS, CEF, LEEF, FortiSIEM EDM, and Google UDM.

New Microsoft Sentinel ASIM reference skill for querying and mapping events to the Azure Sentinel schema.

New Splunk CIM reference skill covering the Common Information Model field mappings.

New Elastic Common Schema (ECS) reference skill for ECS field definitions and event categorization.

New ArcSight CEF reference skill for working with Common Event Format log structure.

New IBM QRadar LEEF reference skill covering Log Event Extended Format field mappings.

+3 more

WAS THIS USEFUL? 👍 👎

Splunk Security Content

Feedback ⓘ v6.1.0 NOTES ↗

Splunk Security Content

v6.1.0 adds 34 new detections and 5 analytic stories covering Cisco IOS XE, SD-WAN, Linux kernel abuse, BlueHammer/RedSun, and PTC Windchill exploitation.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v6.1.0
https://github.com/splunk/security_content.git
# already have the repo? check out this version:
$ git checkout v6.1.0
```

Adds two Cisco Secure Access detections identifying outbound access to anonymizer/privacy services and automated web reconnaissance via abnormal HTTP error patterns.

Adds new analytic stories and detections for BlueHammer and RedSun exploit families targeting Microsoft Defender for privilege escalation and credential access, including suspicious Defender engine/signature updates, unauthorized file modifications, non-admin password changes, and abnormal password reset bursts.

Introduces support for Windows Security Event ID 4723 (password change attempts) to improve visibility into credential theft and privilege escalation tradecraft.

Expands Linux detection coverage for CVE-2026-31431 (Copy Fail) with new analytics targeting malformed authentication entries, `PF_ALG` registration outside boot windows, suspicious namespace creation, and process execution with null argv values.

Introduces a new Linux kern.log data source, providing deeper visibility into low-level system activity and emerging Linux kernel exploitation techniques.

+6 more

└──▷ BREAKING ON UPGRADE

! The detections Attempt To Add Certificate To Untrusted Store, CHCP Command Execution, Ivanti Sentry Authentication Bypass, Processes launching netsh, and Sc exe Manipulating Windows Services have been permanently removed in v6.1.0 with no replacement content; any alerting or workflows depending on these detections will stop functioning after upgrading.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

Triage indicators faster with VirusTotal verdicts in Feedly (2026-07-07) BLOG ↗

Feedly Threat Intelligence is a platform that aggregates and analyzes cybersecurity threat data from multiple sources to help security teams prioritize and respond to emerging threats.

Feedly now embeds VirusTotal verdicts, WHOIS, certs, and reputation scores directly inside IoC Insights Cards.

└──▷ HOW TO FIND IT

Activate VirusTotal enrichment so detection verdicts appear on every IoC Insights Card your team reviews.

> 📍 In Feedly, go to Account Settings, locate the VirusTotal integration section, and paste your existing VirusTotal API key to enable enrichment across all IoC Insights Cards.

Adds VirusTotal integration: connect your existing VirusTotal API key to surface detection verdicts and reputation scores inside IoC Insights Cards.

Embeds WHOIS records, certificate data, network attribution, and community votes on IoC Insights Cards for zero-detection domain analysis.

Enables one-click pivot from an IoC Insights Card directly to the full VirusTotal dataset for deeper investigation.

VirusTotal enrichment data remains account-scoped — never stored by Feedly and never shared with other users or teammates.

WAS THIS USEFUL? 👍 👎

IntelOwl

Feedback ⓘ v6.7.0 NOTES ↗

IntelOwl: manage your Threat Intelligence at scale

IntelOwl v6.7.0 ships an LLM-powered threat-intel chatbot, two new analyzers, and connector health checks.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v6.7.0
https://github.com/intelowlproject/IntelOwl.git
# already have the repo? check out this version:
$ git checkout v6.7.0
```

└──▷ **TRY IT**

Spin up IntelOwl with the local Ollama LLM service enabled to activate the chatbot for your team.

```
$ docker compose -f docker/default.yml -f
docker/optional_analyzers.yml -f docker/override-ollama.yml up
```

Trigger the chatbot's guardrail confirm endpoint via API to approve a pending analyze_observable action proposed by the agent.

```
$ curl -X POST https://<intelowl-host>/api/chatbot/sessions/<session-
id>/confirm -H 'Authorization: Token <your-token>' -H 'Content-Type:
application/json' -d '{}'
```

Check the chatbot health endpoint to verify Ollama is reachable before enabling the chat panel in the UI.

```
$ curl -s https://<intelowl-host>/api/chatbot/health -H
'Authorization: Token <your-token>'
```

Adds a self-deployed LLM chatbot (powered by Ollama via optional Docker override) that lets analysts query IntelOwl jobs, investigations, analyzers, and playbooks in natural language over WebSocket with token streaming.

New `CVE_E` xploitability analyzer combines CISA KEV and FIRST EPSS data to score CVE observables for exploitability.

New RDAP analyzer enables registration data access protocol lookups for domain and IP observables.

Chatbot agent supports an `analyze_observable` action tool with a preview-only guardrail and a separate confirm endpoint, requiring explicit user approval before triggering new analysis jobs.

Chatbot enforces per-user rate limiting on both REST and WebSocket entry points.

+7 more

WAS THIS USEFUL? 👍 👎

Flashpoint

Feedback ⓘ

June 25, 2026: Introducing the Flashpoint MCP Server (2026-06-25) BLOG ↗

Flashpoint is a threat intelligence platform that collects and analyzes data from deep, dark, and open web sources to help organizations identify and mitigate security risks.

Flashpoint MCP Server is now GA, exposing threat intelligence to AI assistants via the Model Context Protocol.

New Flashpoint MCP Server enables AI assistants (e.g., Claude, Gemini) to query Flashpoint intelligence directly via the Model Context Protocol (MCP) standard.

Supports natural-language retrieval across Communities Intelligence, Finished Intelligence Reports, News, Technical Indicators, Vulnerability Intelligence, Ransomware Activity, Compromised Credentials (Enterprise & Customers), Infected Hosts, and Strategic Entity Data (SED).

Authenticates via OAuth or API token, automatically inheriting the organization's existing Flashpoint subscription entitlements.

Read-only access model: AI systems can search, retrieve, enrich, and analyze intelligence without the ability to alter settings, create alerts, or modify operational workflows.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ Interact with apps directly from a case (2026-07-07)

4 RELEASES · 2026-06-23 → 2026-07-07 BLOG ↗ LATEST STABLE

Tines is a no-code automation platform that enables security and operations teams to build powerful workflows that automate repetitive tasks and processes.

Tines now lets you trigger app actions directly from within a case, without leaving the case view.

└──▷ HOW TO FIND IT

Trigger an app (e.g., send a Slack message or create a Jira ticket) without leaving an active case.

> 📍 In the case view, open the Case Actions menu and select the app you want to interact with.

Enables running apps directly from a case via case actions, keeping response workflows in context.

▾ 3 more releases in this issue · 2026-06-23 → 2026-07-07

Filter and sort by Record ID (2026-07-06) BLOG ↗ STABLE

Tines adds Record ID filtering and sorting to records tables and charts.

└──▷ HOW TO FIND IT

Quickly locate a specific record or range of records by ID when triaging a large records table.

> 📍 In the Tines console, open your Records table or chart, select the Filter option, choose `Record ID` as the filter field, enter your target ID or range, then apply Sort by Record ID to order results.

User read and manage tenant permissions (2026-07-01) BLOG ↗ STABLE

Tines adds two new tenant permissions for finer-grained user management control.

Query and aggregate records from the storyboard (2026-06-23) BLOG ↗ STABLE

Tines storyboard record action gains a Query mode for running aggregations directly on record types.

└──▷ HOW TO FIND IT

Aggregate record data inline during a workflow — e.g. count open alerts by severity — without a separate lookup step.

> 📍 In the storyboard, add or open a Record action, select the target record type, then switch the mode to `Query` and configure your aggregation expression.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Tenable Nessus

Feedback ⓘ Tenable Nessus10.12.0 (2026-04-23)

2 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Tenable Nessus is a vulnerability scanner that identifies security weaknesses, misconfigurations, and compliance issues across networks, systems, and applications.

Nessus 10.12.0 adds drag-and-drop UI, Sensor Proxy support, FIPS-140.3, OpenSSL 3.5, ARM64, and a new `--scan-zones` CLI parameter.

└──▷ TRY IT

Link a managed Nessus scanner to Tenable Security Center and assign it to specific scan zones at link time — useful when deploying scanners into segmented network zones.

```
$ nessuscli link --host=<sc_host> --key=<link_key> --scan-zones=
<zone_id>
```

Deploy a Nessus container and assign scan zones via environment variable so the scanner registers to the correct Security Center zone without manual post-link configuration.

```
$ docker run -e SCAN_ZONES=<zone_id> -e LINKING_KEY=<link_key> -e
SC_HOST=<sc_host> tenable/nessus
```

Import a .nessus results file into Nessus without using the file picker — drag the file from your desktop and drop it anywhere in the Nessus UI.

> 📍 In the Nessus UI, drag a .nessus file from your local file manager and drop it directly onto the Nessus scans page to import it.

Adds drag-and-drop scan organization: move existing scans into folders directly from the list view in the modernized UI.

Enables drag-and-drop .nessus file import: drop scan files from your local machine directly into the Nessus UI.

Adds support for OpenSSL 3.5.

Adds support for FIPS-140.3 compliance mode.

Enables Sensor Proxy connectivity from Nessus to Tenable Security Center 6.9.0 and later.

+4 more

└──▷ **BREAKING ON UPGRADE**

! Unauthenticated requests to download files from Tenable Nessus are no longer permitted; those requests must now be authenticated with a session token.

! API access to plugin rules has been removed for users with Basic permissions.

▾ **1 more release in this issue** · seen 2026-07-13

> Tenable Nessus10.12.1 (2026-06-25) NOTES ↗ STABLE
>
> Nessus 10.12.1 adds variable support for custom audit files when scanning from Tenable Security Center.

WAS THIS USEFUL?  

Brinqa

Feedback ⓘ 11.33.98 6 RELEASES · SEEN 2026-07-13 NOTES ↗

LATEST STABLE

Brinqa is a cyber risk management platform that aggregates, correlates, and prioritizes security findings across an organization's attack surface to drive remediation workflows.

Brinqa 11.33.98 adds drill-down support on dashboard charts with complex BQL relationship queries.

Adds drill-down functionality on dashboard charts that use complex BQL queries with relationship conditions (e.g., `FIND Asset THAT HAS Finding`), enabling interactive exploration of data previously locked behind static charts.

Extends due date behavior now preserves the extended due date on findings tied to expired exception requests, preventing retroactive out-of-SLA marking for findings remediated within their approved extension window.

▾ 5 more releases in this issue · seen 2026-07-13

12.0.248 NOTES ↗ STABLE

Brinqa 12.0.248 adds case-insensitive search and scalability improvements for high-volume environments.

Adds case-insensitive search across the platform.

Enhanced platform scalability and performance capabilities to support high-volume customer requirements.

Improved orchestration performance to meet high-volume customer requirements.

Adds validation and sanitization for invalid date/time values during data integration sync.

└──▷ BREAKING ON UPGRADE

! GraphQL and GraphQL Explorer have been removed in version 12.

! Indicator functionality is not supported in version 12.

! Analytics Source functionality is not supported in version 12.0; dashboards that relied on Analytics Source are not available and may require reconfiguration.

! Brinqa Agents must be compatible with version 12; when configuring integrations, Cloud Server must be selected from the Server dropdown.

12.0.321 NOTES ↗ STABLE

Brinqa 12.0.321 adds Source Data Computation automation, new orchestration stages, and expanded flow step timeouts.

Introduces Source Data Computation to automatically identify and process Source Data Models (SDM) with calculated attributes.

Adds new orchestration stages enabling automations to trigger after source data computation and data lifecycle management (DLM) pipeline steps.

Increases the data lifecycle flow step timeout from 1 to 5 minutes, accommodating larger dataset processing.

12.1.118 NOTES ↗ STABLE

Brinqa 12.1.118 adds delta computes, consolidation, and Analytics Source support.

Adds delta computes and consolidation (feature-flagged), enabling incremental data processing workflows.

Adds Analytics Source support, expanding the range of data sources available for ingestion and analysis.

Adds BQL Optimizer feedback loop to improve query planning over time.

Surfaces invalid attribute type mismatches during consolidation in both logs and the UI for improved consolidation visibility.

12.1.403 NOTES ↗ STABLE

Brinqa 12.1.403 adds BCL-based connector filtering, multi-hop BQL traversal, bulk activate/deactivate actions, and richer API pagination.

└──▷ **HOW TO FIND IT**

Bulk-deactivate outdated risk scoring models at once instead of toggling each individually.

> 📍 Navigate to Configuration › Risk Scoring Models, select the target models using the checkboxes, then choose Deactivate from the bulk actions menu.

Adds BCL condition-based filtering during connector sync, letting users control which records are imported based on attribute values.

Adds nested projection and wildcard field selection to the BQL API `returningFields` parameter for precise control over query response attributes.

Adds row count details to BQL API pagination responses, simplifying the construction of pagination-aware integrations.

Adds multi-hop relationship traversal in BQL, enabling recursive queries across hierarchical data such as multi-level organizational structures.

Adds Activate and Deactivate bulk actions to configuration list views for risk factors, SLA definitions, risk scoring models, and status configurations.

12.2.201 NOTES ↗ STABLE

Brinqa 12.2.201 debuts SmartFlows v1.0, an AI Attribution Engine, BrinqaIQ enhancements, and a new Platform API replacing GraphQL.

Introduces SmartFlows v1.0, a new workflow automation capability for building and executing security flows.

Adds ConnectIQ: AI Attribution Engine v1.0 for AI-driven attribution of security findings.

Enhances BrinqaIQ Assistant with expanded AI-driven capabilities.

Replaces the GraphQL API with a new Platform API — a breaking change requiring migration of existing integrations.

Adds a Configuration Change Management UI for building and managing flows visually.

+2 more

└─▷ **BREAKING ON UPGRADE**

! GraphQL API is replaced with Platform API

WAS THIS USEFUL? 👍 👎

Edgescan

Feedback ⓘ

Introduction to Shared Credentials, Browser Recordings and MFA in Edgescan

14 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Edgescan is a cloud-based vulnerability management platform that continuously discovers, assesses, and prioritizes security risks across web applications, APIs, and network infrastructure.

Edgescan adds Shared Credentials, Browser Recordings, and MFA-aware scanning workflows (Beta, GA Q2 2026).

└──▷ **HOW TO FIND IT**

Configure a reusable TOTP-enabled credential and attach it to a browser recording for fully automated MFA login during scans.

📍 In the console, go to Shared Credentials › Create Credential, enter username and password, enable MFA, upload or scan the OTP QR code to populate the OTP URI, save the credential. Then go to Browser Recordings › Create Recording, capture the login flow including the OTP input step, and map the otp_code field from the Shared Credential to the OTP input step. Finally, assign both the Shared Credential and the Browser Recording to the target asset under its Authentication Workflow settings.

Push an out-of-band OTP code to Edgescan via API so an API-push MFA workflow can complete authentication during an assessment.

```
$ curl -X POST https://<edgescan-
instance>/api/v1/mfa/<mfa_record_id>/otp \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"otp_code": "<one_time_code>"}'
```

Introduces a central Shared Credentials library at the organisation level, letting practitioners create one credential set and assign it to any number of assets instead of re-entering credentials per asset.

Adds Browser Recordings that capture real browser login flows — including multi-page, JavaScript-heavy, and MFA-prompted flows — and replay them during authenticated assessments.

Supports mapping Shared Credential fields (username, password, otp_code) to specific input steps in a browser recording, enabling credential reuse across multiple assets with the same recording.

Enables MFA-aware credentials with support for TOTP/authenticator apps, email-based OTP, SMS-based OTP, and API-push/out-of-band MFA patterns.

Supports QR code upload or scan during Shared Credential setup to automatically populate OTP URI and related TOTP fields.

+4 more

▾ **13 more releases in this issue** · seen 2026-07-13

Release Notes - 11th Mar 2026 NOTES ↗ STABLE

Edgescan adds vulnerability detection history timeline, visible discard notes in reports, and expanded MFA/credential workflow capabilities.

└──▷ **HOW TO FIND IT**

Review every re-detection event for a vulnerability during a retesting cycle to build an audit trail.

📍 In the Edgescan platform, open a vulnerability detail page and select the Vulnerability detection history timeline tab to see each re-detection entry logged by assessments.

Ensure pentest due dates are tracked at asset creation time so teams can prioritise PT scheduling.

📍 In the Edgescan platform, go to Assets › Create Asset, complete the asset details, and populate the PT to be completed by date field before saving (Alpha/Beta feature).

Speed up MFA credential setup by scanning an authenticator QR code directly from a webcam instead of manual entry.

📍 In the Edgescan platform, go to Shared Credentials › Create/Edit Credential, select OTP MFA, and choose Scan QR Code to capture the QR code via webcam and auto-populate the OTP secret (Alpha/Beta feature).

Adds a Vulnerability detection history timeline view that logs a new entry each time a vulnerability is re-detected during an assessment, providing richer retest audit trails.

Surfaces discard notes directly in the platform UI and in reports, and allows updating discard notes on already-discarded vulnerabilities.

Enables event notifications for changes to the superuser notes field.

Introduces a PT to be completed by due-date field in the asset creation workflow for pentest targets (Alpha/Beta).

Adds post-authentication stabilization options configurable within authentication workflows (Alpha/Beta).

+3 more

Release Notes - 14th Apr 2026 NOTES ↗ STABLE

Edgescan adds checkbox-based bulk discard for vulnerability management alongside existing filter-based bulk actions.

└──▷ HOW TO FIND IT

Discard a specific subset of findings by hand-picking them with checkboxes rather than relying solely on filter criteria.

📍 In the Vulnerabilities view, select individual findings using the checkboxes, then choose Bulk Discard from the actions menu to discard only the checked items.

Adds checkbox-based bulk discard option, letting users selectively discard individual findings alongside the existing filter-matched bulk-discard workflow.

Release Notes - 14th Jan 2026 NOTES ↗ STABLE

Edgescan adds recorded login sequence uploads, single insight category views, and dark mode for insight graphs.

└──▷ HOW TO FIND IT

When configuring authenticated scans, upload a recorded login sequence so Edgescan can replay credentials automatically — useful for apps with complex SSO or multi-step login flows.

📍 In the platform, go to Configure Scanning › Credentials, record your login sequence using the browser recorder, then upload the recording file to associate it with your asset's scan workflow.

When presenting to a CISO, isolate a single insight category to focus the dashboard view without noise from other report types.

📍 In the platform, navigate to Insights and select a single category (e.g., CISO Report) from the category filter to view only that insight group.

Enables users to record a login sequence and upload it to Edgescan for automated authentication during scanning.

Supports viewing a single insight category in isolation (e.g., CISO Report) within the Insights dashboard.

Adds dark mode support for insight graphs in the platform UI.

Release Notes - 14th May 2026 NOTES ↗ STABLE

Edgescan adds saveable Asset Export filters, duplicate credential detection, in-asset credential creation, and bulk workflow credential vault support.

└──▷ **HOW TO FIND IT**

When managing many assets at scale, supply credential vault IDs during a bulk asset create so scanners authenticate without manual per-asset configuration.

> 📍 In the platform, go to Assets › Bulk Create/Update, and include the Credential Vault ID column in your CSV upload to associate vault credentials with each asset in one operation.

When auditing your credentials vault for redundancy, use the new duplicate filter to surface and consolidate repeated credentials before they cause scan conflicts.

> 📍 In the platform, go to Credentials Vault, apply the Duplicates filter to identify matching credentials, then use the compare and merge controls to consolidate them.

When setting up a new asset scan, create and apply credentials inline without leaving the asset configuration view.

> 📍 In the platform, open an Asset › Scan Configuration, select `Create Credential` within the credentials section, fill in the credential details, and apply it directly to the asset scan.

Saves export column selections alongside saved Asset Export filters, eliminating the need to reconfigure columns when switching between filters.

Adds controls to identify, filter, compare, and merge duplicate credentials directly within the credentials vault.

Enables credential creation and application from the asset scan configuration view, removing the need to navigate away from the asset page.

Updates the AI Insight Licence Forecast Report with org-level visibility into licence usage, expiries, and growth forecasts for capacity planning and renewal management.

Adds a legend to complex Mermaid line graphs in Insights for easier series identification.

+1 more

Release Notes - 15th June 2026 NOTES ↗ STABLE

Edgescan adds asset assessment date filtering, AI Insights UI refresh, scoped token mappings, and a raft of workflow enhancements.

└──▷ **HOW TO FIND IT**

Configure a custom scope directive for a token mapping to restrict a session token to a specific path, preventing it from leaking to out-of-scope endpoints.

> 📍 In the console, go to the asset's authentication settings, open Token Mappings, and add a custom Scope Directive tied to the relevant location specifier for the token.

Filter vulnerabilities across multiple severity levels in a single view to prioritise remediation across Critical and High findings together.

> 📍 In the Vulnerabilities list, open the severity filter and select multiple severity levels (e.g. Critical and High) to apply a combined filter.

Scope a scheduled assessments export to a single business unit and its child organizations for targeted reporting.

> 📍 In Scheduled Exports, create or edit an assessments export and set the Organization (with children) scope to the target organization before saving.

Adds `Date of last assessment` column and filter, letting users view and filter assets by the date of a specific assessment type (e.g. last pen-test date, last retest date).

Refreshes the AI Insights page with a look and feel consistent with the rest of the Edgescan interface.

Enables scoped scheduled assessment exports to a single organization and its descendants.

Adds multi-severity filtering so users can filter vulnerabilities by multiple severity levels simultaneously.

Adds retest-in-progress status to the asset badge (e.g. `scanning`, `validating`).

`+8 more`

`Release Notes - 16th Oct 2025` **NOTES ↗** **STABLE**

Edgescan adds tester notes for retests, vulnerability definitions in reports, tag visibility, and new vuln filters.

└──▷ HOW TO FIND IT

Include inline vulnerability definitions in a report so reviewers don't need to cross-reference external documentation.

> 📍 In the console, go to Reports › Generate Report, then enable the `Definitions` option to embed each vulnerability's definition directly after its entry in the output.

Scope vulnerability triage by filtering out out-of-scope findings or narrowing by asset status on the Vulnerabilities page.

📍 In the console, go to Vulnerabilities, open the filter panel, and apply the `Out of Scope` filter and/or the `Asset Status` filter to narrow results.

Communicate context to the operations team before a retest by attaching notes at the assessment or vulnerability level.

📍 In the console, when initiating a retest assessment, enter guidance in the `Tester Notes` field under assessment or vulnerability notes before submitting.

Auto-populates the current asset when navigating to the Report section from an asset details page.

Adds a definitions option in report generation to display each vulnerability's definition inline, immediately after the vulnerability.

Adds a Tags column to the asset table for at-a-glance tag visibility.

Adds new vulnerability filter options: filter by out-of-scope status and by asset status.

Adds a tester notes field allowing users to provide notes to the operations team when initiating a retest assessment, via assessment and vulnerability notes.

`Release Notes - 24th Nov 2025` `NOTES ↗` `STABLE`

Edgescan adds audit trail attribution, bulk asset report selection, and role-based access control.

└──▷ **HOW TO FIND IT**

Review who initiated a retest on a specific vulnerability to support change-management and accountability workflows.

📍 In the console, navigate to the Audit Trail and filter by event type `Retest` or `Vulnerability Test` to see the triggering user alongside each assessment entry.

Generate a consolidated report across a subset of assets in one action rather than running separate per-asset reports.

📍 In the console, go to Reports › Generate Report, use the bulk select checkbox to choose multiple assets, then proceed to generate the report.

Restrict a team's access to only the assets and data relevant to their function by assigning a custom role.

📍 In the console, go to Settings › Roles, create or select a role with the desired resource permissions, then assign it to the target users under User Management.

Adds audit trail entries that identify which user triggered a retest or vulnerability assessment.

Enables bulk selection of multiple assets when generating reports.

Introduces role-based authorization support, allowing created and assigned roles to control access to resources within an Edgescan account.

Release Notes - 28th May 2026 NOTES ↗ STABLE

Edgescan adds asset-creation notifications, AI Insights asset filtering, configurable asset exports, and improved CWE/assessment-export filtering.

└──▷ **HOW TO FIND IT**

Receive an alert in your notification pipeline the moment a new asset is added to Edgescan — useful for triggering automated onboarding workflows.

> 📍 In the Edgescan console, go to Notifications › Events, create a new event, and select the `Asset Created` event type as the trigger condition.

Generate a focused AI Insight scoped to a specific asset group (by ID or tag) without overwriting the shared organization Insight.

> 📍 In the Edgescan console, go to AI Insights, apply the Asset IDs or Tags filter to scope the data, then select `Generate Personal Insight` to produce an individual report.

Produce a clean asset report without risk-over-time graphs — useful for executive summaries where trend charts add clutter.

> 📍 In the Edgescan console, go to Asset Reporting › Export, open the report configuration options, and uncheck `Include risk-over-time graphs` under the Asset List section before exporting.

Adds a new notification event triggered when an asset is created, enabling teams to react in real time to asset inventory changes.

Enables filtering of AI Insights data by asset IDs and tags, and supports generating personal Insights without affecting the organization-wide Insight.

Adds an option to omit risk-over-time graphs from the Asset List section of asset reports for cleaner exports.

Supports both numeric (e.g. 79) and full string (e.g. CWE-79) input when filtering vulnerabilities by CWE in the UI.

Applies active UI filters to Assessment Exports so the exported data matches exactly what is currently displayed.

+1 more

Release Notes - 29th Apr 2026 NOTES ↗ STABLE

Edgescan adds webhook note payloads and new CISO AI Insight reporting period and time-unit controls.

└──▷ **HOW TO FIND IT**

Capture full note text in a webhook consumer to automatically route triage notes into a ticketing system when an analyst adds a note in Edgescan.

> 📍 In the console, go to Integrations › Webhooks, select your webhook, and enable the note-added trigger — the event payload will now include the note content field.

Scope a CISO AI Insight report to the last completed quarter for board-ready quarterly risk summaries.

> 📍 In the console, go to Insights › CISO Report, open the reporting period selector, and choose `Last completed quarter`.

Adds note content to webhook event payloads when the note-added trigger fires, enabling downstream systems to consume full note context.

Enables setting the CISO AI Insight report's time period to `last completed quarter` and similar discrete reporting intervals.

Displays FTE time-saved calculations in hours (instead of minutes) on the CISO AI Insight report for clearer workforce-impact communication.

Release Notes - 29th Jan 2026 NOTES ↗ STABLE

Edgescan reports now surface EXF, EPSS, and CISA KEV scoring metrics alongside each vulnerability.

└──▷ **HOW TO FIND IT**

Generate a vulnerability report to see exploit feasibility (EXF), EPSS probability, and CISA KEV status in a single export for faster triage.

> 📍 In the Edgescan console, navigate to Reports, generate a vulnerability report, and review the EXF, EPSS, and CISA KEV columns now included for each finding.

Adds EXF, EPSS, and CISA KEV indicators to generated vulnerability reports for richer prioritization context.

Release Notes - 30th Mar 2026 NOTES ↗ STABLE

Edgescan adds UX improvement to Roles permissions dropdowns, keeping selections open for faster multi-item configuration.

Improves Roles UI so permissions and resources dropdowns stay open after adding items, enabling faster multi-selection without repeated reopening.

Release Notes - 6th Nov 2025 NOTES ↗ STABLE

Edgescan adds Out of Scope discard reason, discard reason in exports, and an automatic retest queue.

Adds Out of Scope as a selectable discard reason when dismissing vulnerabilities.

Includes discard reason field in vulnerability exports for audit and reporting workflows.

New retest queue automatically schedules vulnerability retest requests that cannot fit into the current ongoing retest, eliminating manual retry loops.

Release Notes - 8th Dec 2025 NOTES ↗ STABLE

Edgescan adds pentesting checklists in reports and retest/on-demand assessment start notifications.

Adds a penetration testing checklist to generated reports, summarising what was tested during an engagement.

Sends assessment start notifications when a user invokes a retest or on-demand assessment.

Includes a note in email notifications warning non-registered recipients that report links will not work for them.

WAS THIS USEFUL?  

Cloudmersive

Feedback ⓘ

Cloudmersive Cloud Virus API v11-9-1 (2026-07-11)

3 RELEASES · 2026-06-24 → 2026-07-11 BLOG ↗ LATEST STABLE

Cloudmersive Release Notes is an RSS feed that tracks and publishes updates, improvements, and changes made to Cloudmersive's suite of APIs and developer tools.

Cloudmersive Cloud Virus API v11-9-1 adds enhancements to the Scan Website capability.

Enhances Scan Website functionality with new improvements for web-based threat detection.

▾ 2 more releases in this issue · 2026-06-24 → 2026-07-11

> Cloudmersive Cloud Virus API v11-8-8 (2026-07-04) BLOG ↗ STABLE
>
> Cloudmersive Cloud Virus API v11-8-8 adds detection enhancements to its virus scanning service.

> Cloudmersive Cloud Virus API v11-7-3 (2026-06-24) BLOG ↗ STABLE
>
> Cloudmersive Cloud Virus API v11-7-3 adds virus scanning support for .vpd and .fsdb file formats.
>
> Supports scanning of .vpd and .fsdb file formats for viruses and malware.

WAS THIS USEFUL? 👍 👎

ssh-audit

Feedback ⓘ v3.9.0 NOTES ↗ CODE ↗

SSH server & client security auditing (banner, key exchange, encryption, mac, compression, compatibility, security, etc)

ssh-audit v3.9.0 adds post-quantum warnings, UNIX socket scanning, SOCKS5 proxy support, and built-in hardening guides.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.9.0 https://github.com/jtesta/ssh-audit.git
# already have the repo? check out this version:
$ git checkout v3.9.0
```

└──▷ TRY IT

Audit an SSH daemon exposed only via a UNIX socket inside a container or CI/CD pipeline — no TCP port required.

```
$ ssh-audit unix:///run/ssh-unix-local/socket
```

Scan an internal host through a SOCKS5 jump proxy — useful when direct connectivity to the target is unavailable.

```
$ ssh-audit --socks5 proxy.example.com:1080 target-host.internal
```

Retrieve the built-in hardening guide for a specific platform to apply recommended SSH configuration without visiting an external site.

```
$ ssh-audit --list-hardening-guides
ssh-audit --get-hardening-guide "Debian 13"
```

Adds warnings on all key exchanges lacking post-quantum protection, flagging exposure to Harvest Now, Decrypt Later attacks.

New `--list-hardening-guides` and `--get-hardening-guide` flags surface built-in hardening guides for common platforms (Debian, Ubuntu, Rocky Linux, OpenSSH) without requiring internet access.

Supports scanning SSH servers over UNIX sockets (e.g. `unix:///path/to/socket`), enabling auditing in containerized and CI/CD environments.

New `--socks5 host:port` flag enables scanning through a SOCKS5 proxy (implies `--skip-rate-test`).

Identifies 2 new post-quantum key exchanges: `mlkem768nistp256-sha256` and `mlkem1024nistp384-sha384`.

+6 more

└─▷ **BREAKING ON UPGRADE**

- ! Support for Python 3.8 has been dropped, so setups running ssh-audit on Python 3.8 will no longer work.
- ! Support for Python 3.9 has been dropped, so setups running ssh-audit on Python 3.9 will no longer work.
- ! SSHv1 support has been removed, so any scans or configurations relying on SSHv1 scanning will no longer function.
- ! Batch mode (`-b`) no longer automatically enables verbose mode, so scripts or workflows that depended on batch mode implying verbose output will see less output unless `-v` is explicitly added.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Axonius

Feedback ⓘ Axonius Release Notes 9.0.2 (2026-07-09) BLOG ↗

Axonius is a cybersecurity asset management platform that discovers, inventories, and manages all devices, users, and cloud assets across an organization's environment.

Axonius 9.0.2 adds CISA Vulnrichment, Latest Snapshot View, Ingestion Rules set_value, 4 new adapters, and SAML Visiting Data Scopes.

└──▷ HOW TO FIND IT

Run stable asset queries during an active discovery cycle without waiting for it to complete.

> 📍 In the Axonius console, open any asset table, click the View selector at the top, choose Latest Snapshot View, then run your query — results reflect the last completed discovery and remain stable while a new cycle is in progress.

Dynamically normalize vendor name casing during adapter ingestion to keep asset data consistent across sources.

> 📍 In Axonius console, go to Ingestion Rules › create or edit a rule › add a Post-Action of type set_value, set the target field to `vendor`, choose the trim_prefix or regex_extract transformation, and save — the rule applies the transformation to every record ingested from the configured adapter.

Prevent stale resolved CVEs from reappearing in the platform after a re-fetch.

> 📍 In the Axonius console, navigate to an adapter's Advanced Settings › Adapter Configuration, locate 'Ignore Security Findings fixed more than X days ago', set X to your desired retention window (e.g., 30), and save the connection.

Adds CISA Vulnrichment as an out-of-the-box vulnerability enrichment source, providing CWE mappings, affected product data, and SSVC metrics.

New CVE Enrichment Settings section under System Settings > Enrichment consolidates control over 24 vulnerability enrichment sources with granular field selection.

Introduces Latest Snapshot View — a point-in-time, read-only view of asset data isolated from live discovery cycles, enabling faster and stable query results during active discovery.

Rebuilt data tables from the ground up, removing column limits, adding column pinning in Edit Columns, and replacing hover-based row actions with a three-dot menu.

Adds Markdown support to Dashboard Text Widgets, enabling formatted headings, bold text, and other markup directly in the widget editor.

+22 more

└─▷ **BREAKING ON UPGRADE**

- ! Bitbucket Cloud V2 connections require workspace slugs to be added after upgrading, as the `/user/workspaces` endpoint has been deprecated and a new Workspace Slugs connection setting is now required.
- ! Com Laude domain records have been changed from the Devices asset type to Domains & URLs; existing queries that filter Devices by domain-related fields may need to be updated to query URLs instead.
- ! Infoblox DDI A Records (`record:a`) and Shared A Records (`sharedrecord:a`) are now parsed as Domains & URLs instead of Devices; any saved queries, Enforcement Actions, dashboards, reports, or workflows targeting these records as Devices will need to be updated to use Domains & URLs instead.

WAS THIS USEFUL? 👍 👎

Intruder

Feedback ⓘ Push vulnerabilities to Linear (2026-07-09)

4 RELEASES · 2026-06-15 → 2026-07-09 BLOG ↗ LATEST STABLE

Intruder is a cloud-based vulnerability scanner that continuously monitors and identifies security weaknesses in internet-exposed systems, helping businesses reduce their attack surface.

Intruder adds native Linear integration to push vulnerabilities directly into Linear teams, projects, and backlogs.

.

└──▷ HOW TO FIND IT

Connect Linear and configure auto-create rules so critical and high findings flow into your backlog automatically.

📍 In the Intruder console, go to Integrations → Linear, connect your workspace, then set auto-create rules by severity (e.g. Critical and High) and select the target Linear team and project.

New Linear integration: connect your Linear workspace from Intruder settings to push vulnerabilities without leaving the platform.

Supports one-click push of any issue to a chosen Linear team and project.

Enables auto-create rules to automatically raise Linear tickets based on vulnerability severity.

Displays linked Linear tickets inline in the issue list and issue details panel.

Includes test-issue feature to validate integration configuration before going live.

▾ 3 more releases in this issue · 2026-06-15 → 2026-07-09

Connect Docker Hub for container image scanning (2026-07-02) BLOG ↗ STABLE

Intruder adds Docker Hub integration for automated container image vulnerability scanning within the portal.

.

└──▷ HOW TO FIND IT

Connect a Docker Hub account to bring container images into Intruder's vulnerability monitoring alongside existing cloud targets.

📍 In the Intruder portal, go to Discovery → Container Images → click Add registry (top right) → select Docker Hub → enter your Personal Access Token (PAT) to authenticate and begin repository discovery.

Connects multiple personal or organisation Docker Hub accounts via Personal Access Token (PAT) for container image scanning.

Automatically discovers repositories, including tags and last-pushed dates, with on-demand refresh support.

Enables adding a custom tag to create a scan target and trigger an immediate vulnerability scan.

Runs automatic daily rescans and instant rescans when a tag points to a new image version.

`Regex tag matching for container images (2026-06-18)` BLOG ↗ STABLE

Intruder adds regex tag matching for container image targets, enabling accurate scoping across mixed or non-standard tagging conventions.

.

└──▷ **HOW TO FIND IT**

Target only production-tagged images across repos that mix semver and custom suffixes, without manually updating the tag each release.

> 📍 In the console, go to your container image target › Tag rule › select `Regex`, then enter your pattern (e.g. ^v\d+\.\d+\.\d+$) and save.

New regex tag-matching rule for container image targets lets you define a custom pattern to consistently resolve the correct image version.

Supports dynamic targeting: the target automatically follows the latest tag matching the regex as new image versions are pushed.

Enables accurate scoping across repositories with mixed conventions (e.g., semver, `-dev` suffixes, or custom schemes) where exact-match and semver rules fall short.

`Attack Surface view for Cloud and Pro plans (2026-06-15)` BLOG ↗ STABLE

Intruder launches Attack Surface view: unified perimeter visibility with port, service, and web screenshot coverage.

.

└──▷ **HOW TO FIND IT**

Spot every host exposing a database port (e.g. 5432) across your entire perimeter before an auditor or attacker does.

> 📍 In the Intruder portal, go to Attack Surface (left nav) › filter by Port: 5432 to see every target with PostgreSQL exposed.

Identify certificates expiring soon across all targets to prevent unexpected lapses in production services.

📍 In the Intruder portal, go to Attack Surface (left nav) › filter by Certificate Expiry to surface certs about to lapse.

New Attack Surface view consolidates every exposed port across all monitored targets into a single list, showing the service and software version behind each.

Displays screenshots of every exposed web service so practitioners can instantly distinguish forgotten admin panels from intended login pages.

Enables search and filtering by port, service, or product to identify all targets running a specific service in seconds.

Adds filtering by host status or certificate expiry to surface newly appeared services or TLS certificates approaching lapse.

Available on Cloud and Pro plans, with port coverage scaling by plan tier.

WAS THIS USEFUL? 👍 👎

runZero

Feedback ⓘ runZero 5.0.260707.0 (2026-07-07)

7 RELEASES · 2026-06-17 → 2026-07-07 BLOG ↗ LATEST STABLE

runZero is a network discovery and asset inventory platform that scans environments to identify and manage all connected devices, including unmanaged and OT assets.

runZero 5.0.260707.0 adds per-source stale threshold overrides and a software count drill-down link in asset inventory.

└──▷ HOW TO FIND IT

Set a tighter stale threshold for a high-churn integration source (e.g., a cloud connector) without affecting the global expiration policy for all other integrations.

> 📍 In the console, go to Organization Settings (or Account Settings) › Integrations, select the integration source, and set a stale threshold override value for that source.

Supports per-integration stale threshold overrides for the stale integration attribute expiration feature, configurable in organization and account settings; integrations without an override continue to use the global setting.

Software count column in the asset inventory now links directly to the software inventory for faster drill-down.

▾ 6 more releases in this issue · 2026-06-17 → 2026-07-07

> runZero 5.0.260702.0 (2026-07-02) BLOG ↗ STABLE
>
> Dragos integration gains serial-number-based asset matching with SentinelOne; suppressed vulnerabilities now count as Resolved in Issue status.
>
> Dragos integration now matches assets with SentinelOne data using asset serial number (bidirectionally).

> runZero 5.0.260630.0 (2026-06-30) BLOG ↗ STABLE
>
> runZero 5.0.260630.0 adds valueless tag search to asset inventory queries.
>
> └──▷ HOW TO FIND IT
>
> Find assets tagged `test` with no assigned value — useful for auditing incomplete or placeholder tags in the inventory.
>
> > 📍 In the asset inventory search bar, enter: tag:test=
>
> Supports searching for assets with a tag that has no value (e.g., `tag:test=`) to distinguish unvalued tags from valued ones like `tag:test=example`.

runZero 5.0.260625.0 (2026-06-25) BLOG ↗ STABLE

runZero 5.0.260625.0 adds AWS VPC Endpoint asset discovery and external issue sharing.

Adds AWS VPC Endpoint asset creation to the AWS integration, enabling visibility into VPC Endpoint resources.

Enables external sharing of issues directly from the issues datagrid.

runZero 5.0.260625.0 - Detect, Prioritize, Remediate, Verify, Report (2026-06-25)

BLOG ↗ STABLE

runZero 5.0 delivers end-to-end exposure management: remediation tracking, verified closure, recurring reports, and broader OT/IT vulnerability detection.

└──▷ TRY IT

Enable or disable version-based vulnerability matching for a specific organization via the public API to control Missing Patches rollup reporting.

```
$ curl -X PATCH https://<your-runzero-instance>/api/v1.0/org/<org-id> \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vmatch_enabled": true}'
```

Schedule a recurring weekly vulnerability report for external stakeholders who need no runZero account to review findings.

📍 In the console, go to Reports, select a supported view (e.g., Vulnerability Groups), choose Schedule Report, set recurrence to Weekly, add recipient email addresses, and save — recipients receive an interactive HTML snapshot with CSV/JSON export options.

Initiate a remediation workflow and link it to Jira directly from a prioritized exposure finding.

📍 In the console, open an exposure, finding, or vulnerability group, click Create Remediation Issue, confirm the auto-populated Name/Description/Remediation fields, select your Jira credential and project key, configure the state mapping, and click Push as External Ticket.

New Exposure Management dashboard as the default landing view, surfacing prioritized exposures with separation of OT/IoT streams, MTTR tracking, issue status

totals, reopened-issue trends, vulnerability status trends by severity, and a Goals progress widget.

Adds Remediation Workflows: create a Remediation Issue in one click from an exposure, finding, or vulnerability group, with auto-populated fields and push-to-Jira integration including custom state-to-state mapping.

Enables verified remediation closure: runZero preserves a compressed snapshot of the original exposure state and compares it against live discovery on scheduled scans to confirm fixes, with automatic reopened-issue tracking.

Adds searchable, sortable Recurrences and Last Recurrence columns to the Issues inventory, and due-date alert rules that fire email or webhook notifications when a Remediation Issue exceeds a configured number of days past due.

Adds Recurring External Reporting: schedule weekly or monthly interactive HTML snapshot reports from supported dashboards and views, deliverable by email to recipients without a runZero account, with CSV, JSON Lines, JSON, and PNG export options.

+6 more

runZero 5.0.260618.0 (2026-06-18) BLOG ↗ STABLE

runZero 5.0.260618.0 adds all-asset alert queries, on-demand alert runs, expanded GE-SRTP ports, ADB STLS metadata, and richer InsightVM integration.

Adds option in alert rule queries to cover all assets, not just live assets.

Enables on-demand execution of alert rules without waiting for a scheduled run.

Expands scanner coverage to additional default ports for the GE-SRTP protocol.

Captures metadata from ADB STLS service responses during scanning.

Includes asset hostnames in the Rapid7 InsightVM Cloud Integration.

+1 more

runZero 5.0.260617.0 (2026-06-17) BLOG ↗ STABLE

runZero 5.0.260617.0 adds issue type and external ticket link columns to the Issues table.

WAS THIS USEFUL? 👍 👎

Censys

Feedback ⓘ July 6, 2026 (2026-07-06)

4 RELEASES · 2026-06-16 → 2026-07-06 BLOG ↗ LATEST STABLE

Censys is an internet intelligence platform that continuously scans and indexes internet-connected devices, helping security teams discover and monitor exposed assets and vulnerabilities.

Censys Active DNS goes GA, adds Live Resolution, 11 new fingerprints, GOPHER scanner, and 2 new ASM CVE risk fingerprints.

└──▷ **HOW TO FIND IT**

Surface all exposed Cacti servers in your ASM workspace that are vulnerable to the newly tracked pre-auth SQL injection CVE.

> 📍 In the ASM web console, go to Risk Types configuration and enable `Vulnerable Cacti Server [CVE-2026-39955]`, then filter your risks view with: risks.name: `Vulnerable Cacti Server [CVE-2026-39955]`

Check whether a domain your team is investigating currently resolves in DNS, even if it isn't yet indexed in Censys datasets.

> 📍 In Censys Search or Censys Core, open the Active DNS section, enter the target FQDN, and select `Live Resolution` to retrieve real-time DNS records or confirm the domain does not currently resolve.

Hunt for Wordfence Security plugin installations across the internet to identify WordPress assets in your scope.

```
$ In Censys Search, run the query:
services.software.product:`Wordfence Security`
```

Censys Active DNS is now generally available to all Censys Search and Censys Core users, enabling domain-to-host relationship mapping, infrastructure change analysis over time, and extensive DNS record review.

New Live Resolution feature lets users retrieve DNS records for domains not currently present in Censys datasets, with explicit confirmation when an FQDN does not resolve.

Adds 11 new platform fingerprints covering hardware (D-Link DCS Network Camera), firmware (Lantronix G Series), and software (Wordfence Security, Advanced Custom Fields, All in One SEO, Autoptimize, and more WordPress plugins).

New GOPHER protocol scanner added to the Platform, available to Starter, Search, and Core users.

Two new ASM risk fingerprints added for Vulnerable Cacti Server CVE-2026-39938 (unauthenticated local file inclusion / potential RCE) and CVE-2026-39955 (pre-auth SQL injection in graph_view.php) on Cacti 1.2.30 and prior.

▾ **3 more releases in this issue** · 2026-06-16 → 2026-07-06

June 29, 2026 (2026-06-29) BLOG ↗ STABLE

Censys adds 16 new software/hardware fingerprints and a Gogs RCE risk detection (CVE-2026-52813) to Platform and ASM.

└──▷ **TRY IT**

Identify internet-exposed OpenVPN instances in your ASM inventory to audit unexpected VPN endpoints.

```
$ Platform query: software.product: OpenVPN
```

Surface assets in your ASM workspace that are vulnerable to the Gogs RCE (CVE-2026-52813) for immediate patching prioritization.

```
$ risks.name: "Vulnerable Gogs [CVE-2026-52813]"
```

Enable the new Gogs CVE-2026-52813 risk type if it was disabled by default in your workspace.

📍 In the ASM web console, go to Settings › Risk Type Configuration and enable `Vulnerable Gogs [CVE-2026-52813]`.

Adds 15 new Platform software fingerprints covering OpenVPN, Tor, TeamViewer, FreeSWITCH, OpenFlow, RTMP, Redfish, HL7, Bubble, Lovable, Replit, Synergy, EZVIZ IP Camera, GeoVision NVR, Reolink IP Camera, and Windows CE embedded OS.

New ASM risk fingerprint detects hosts running Gogs versions vulnerable to CVE-2026-52813, an unauthenticated path traversal enabling remote code execution via Git hook overwrite.

All new fingerprints are queryable in the Censys Platform; new ASM risk types may be disabled by default and must be reviewed in the ASM web console risk type configuration.

June 22, 2026 (2026-06-22) BLOG ↗ STABLE

Censys adds bulk API asset tagging, Okta SCIM provisioning, 8 new protocol scanners, and 6 new software fingerprints.

└──▷ **TRY IT**

Identify all exposed Splunk instances in your ASM workspace that are potentially vulnerable to the newly tracked CVE-2026-20253 RCE flaw.

```
$ risks.name:"Vulnerable Splunk Enterprise [CVE-2026-20253]"
```

Confirm new ASM risk fingerprints (like the Splunk CVE risk) are enabled — they may be off by default.

> 📍 In the ASM web console, go to Settings › Risk Type Configuration and review newly added risk types to enable them for your workspace.

Adds bulk asset tagging via API: assign or un-assign a tag to every asset matching a CenQL query in a single operation.

Supports SCIM provisioning via Okta, enabling Platform Admins to automatically provision and de-provision Censys users from their IdP.

Adds 8 new protocol/application scanners: `CLICKHOUSE_NATIVE`, `CLICKHOUSE_HTTP`, OPENFLOW, `REOLINK_BAICHUAN`, RTMP, SYNERGY, `WINCE_CERDISP`, and WORDPRESS.

Adds 6 new software fingerprints: REDCap, Joomla, Joomla Content Editor (JCE), WooCommerce, ChromaDB Admin, and Splunk.

Adds new ASM risk fingerprint detecting Vulnerable Splunk Enterprise instances affected by CVE-2026-20253 (unauthenticated arbitrary file creation/truncation leading to RCE), queryable via `risks.name:Vulnerable Splunk Enterprise [CVE-2026-20253]`.

June 15, 2026 (2026-06-16) **BLOG ↗** **STABLE**

Censys adds a lightweight Host Enrichment API built for high-volume SOC automation in SIEM/SOAR pipelines.

New Host Enrichment API delivers a curated, fixed subset of IPv4/IPv6 host data optimized for high-volume, automated lookups without traditional credit consumption.

Supports SIEM and SOAR integrations, enabling SOC teams to operationalize external internet data at scale.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

Feedback ⓘ API key modification

13 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Twingate is a zero-trust network access tool that securely connects users to private resources without requiring a traditional VPN.

Twingate API keys can now be renamed, have permissions modified, and be disabled/re-enabled after creation.

└──▷ HOW TO FIND IT

Temporarily pause an API key during an incident response without rotating credentials — re-enable it once the situation is resolved.

> 📍 In the Twingate Admin Console, go to Settings › API › API Keys, select the key, and toggle Disable/Enable or update its permissions.

Enables editing of existing API keys to rename them, update their permissions, or disable/enable them without generating new keys.

▾ 12 more releases in this issue · seen 2026-07-13

> Authentication Audit Logs NOTES ↗ STABLE
>
> Twingate adds Authentication Audit Logs with exportable success/failure events for logins, resource access, and MFA.
>
> New Authentication Audit Logs surface per-event success and failure data for Client logins, Resource access attempts, and MFA setup.
>
> Enables export and sync of authentication event data for SIEM ingestion, compliance auditing, and troubleshooting workflows.

> Billing Admin Role NOTES ↗ STABLE
>
> Twingate adds a Billing Admin role that scopes users to plan and payment management only.
>
> └──▷ HOW TO FIND IT
>
> Restrict a finance team member to billing management only, keeping them out of network and access settings.

> 📍 In the Admin Console, go to Team › Users, select the user, and assign the Billing Admin role under their role settings.

New Billing Admin role lets admins grant billing-only access to users who need to manage plans and payments without exposing the rest of the Admin Console.

Connector Deployment Page Improvements NOTES ↗ STABLE

Twingate Admin Console adds Connector deployment guides for Ubiquiti, ZimaOS, and CasaOS.

.

└──▷ **HOW TO FIND IT**

Deploy a Twingate Connector into a ZimaOS or CasaOS homelab environment using the new guided workflow.

> 📍 In the Admin Console, go to the Connector Deployment page, select your target environment (Ubiquiti, ZimaOS, or CasaOS), and follow the step-by-step deployment guide.

Adds step-by-step Connector deployment guides for Ubiquiti, ZimaOS, and CasaOS environments in the Admin Console.

Improves Connector Deployment page UX to make it easier to find the right deployment method for your environment.

Connector Health Metrics NOTES ↗ STABLE

Twingate Connectors now export real-time health metrics in Prometheus format for integration with existing monitoring stacks.

New Connector Health Metrics provide admins visibility into network traffic patterns and connection efficiency via real-time data.

Granular Admin Console search NOTES ↗ STABLE

Twingate Admin Console gains multi-term and exact-phrase search across users, Resources, and devices.

.

└──▷ **HOW TO FIND IT**

Quickly locate a specific resource by its full name when a broad search returns too many results.

> 📍 In the Admin Console, navigate to Resources and enter `Grafana Dashboard` (with double quotes) in the search bar to return only exact matches.

Supports multi-term search across users, Resources, and devices in the Admin Console.

Enables exact-phrase matching by wrapping search terms in double quotation marks (e.g., `Grafana Dashboard`).

`JIT Access Requests` `NOTES ↗` `STABLE`

Twingate adds Just-in-Time access requests with time-limited grants and optional admin approval workflows.

.

└──▷ **HOW TO FIND IT**

Enable a JIT workflow requiring admin approval for a sensitive Resource — ideal for production databases or critical infrastructure in incident response scenarios.

> 📍 In the Admin Console, go to the target Resource › Access › Enable Just-in-Time Access, set approval mode to `Require Admin Approval`, and configure the maximum access duration (e.g., 4 hours).

Configure auto-approved JIT access for a Resource where access needs to be audited but not gated — useful for low-risk Resources where logging is the primary goal.

> 📍 In the Admin Console, go to the target Resource › Access › Enable Just-in-Time Access, set approval mode to `Auto-Approve`, and set the maximum access duration (e.g., 1 day).

New JIT Access Request workflows let admins grant users on-demand, time-limited access to Resources (1 hour to 7 days).

Supports two approval modes: require explicit admin approval or auto-approve for audit-only tracking.

Enables `break glass` incident response protocols by controlling and logging access to sensitive Resources.

`Multi-Account Support` `NOTES ↗` `STABLE`

Twingate Clients now support simultaneous sign-in to multiple accounts across different networks on macOS, iOS, and Windows.

.

└──▷ **HOW TO FIND IT**

Switch between a personal and a corporate Twingate account simultaneously without logging out — useful for admins managing multiple tenant networks.

> 📍 In the Twingate client (macOS, iOS, or Windows), open the account menu and select `Add Account`, then sign in with a second set of credentials. Both accounts

appear in the client and can be individually paused or activated.

Enables simultaneous sign-in to multiple Twingate accounts (e.g., home, work, or multiple admin accounts) from a single client session.

Supports per-account connection pausing without requiring logout or quitting the client.

Multi-account support available on macOS, iOS, and Windows clients, with Android and Linux coming soon.

Network Health Summaries NOTES ↗ STABLE

Twingate adds downloadable Network Health Summary reports with per-user metrics for proactive monitoring and troubleshooting.

.

└──▷ **HOW TO FIND IT**

Download a network health summary to review error rates and data transfer volumes across your Twingate network for a given period.

📍 In the Admin Console, go to Reports › Network Health Summaries, select your desired timeframe, and click Download to export the report.

New downloadable Network Health Summary reports surface metrics such as data transferred and errors across configurable timeframes.

Extends existing active user reports with per-user network health data points, enabling admin-level visibility into individual user connectivity.

Privileged Access for SSH NOTES ↗ STABLE

Twingate adds Privileged Access for SSH: register servers as SSH Resources with IdP auth, cert-based authN, and session recording.

.

└──▷ **HOW TO FIND IT**

Enable SSH Resource session recording output for forwarding to a SIEM — navigate in the Twingate Admin Console to configure the SSH Resource.

📍 In the Admin Console, go to Resources › Add Resource, select SSH as the resource type, specify your server address, and enable session recording to capture asciicast output for SIEM forwarding.

Adds SSH Resources: register servers and VMs in Twingate and let users connect via existing SSH tools without managing or distributing SSH keys.

Enables identity-provider authentication for SSH sessions through the Twingate Identity Firewall, eliminating credential sprawl and bastion hosts.

Gateway handles certificate-based authentication to target servers on behalf of users, removing the need for per-user SSH key management.

Records every SSH session to stdout in asciicast format, ready to forward to a SIEM or replay for audit review.

Resource Approvers NOTES ↗ STABLE

Twingate adds per-Resource Approver Groups so non-Admin teams can review and approve JIT access requests.

.

└──▷ HOW TO FIND IT

Delegate JIT access approvals for a sensitive resource to its owning team so Admins aren't a bottleneck.

📍 In the Admin Console, go to Resources › select a Resource › Access › Resource Approvers, then add one or more Groups whose members should review and approve Access Requests for that Resource.

Enables delegation of JIT Access Request approvals to designated Groups on a per-Resource basis, without granting full Admin privileges.

Supports configuring one or more Groups as Resource Approvers, letting resource-owning teams manage their own access requests.

Resource Exclusions NOTES ↗ STABLE

Twingate adds Resource Exclusions to route specific addresses around Twingate infrastructure while keeping audit-log traceability.

.

Enables per-address bypass routing so specific addresses within a broader Resource scope (e.g., a public subdomain under a wildcarded corporate domain) go directly via the local OS, skipping Twingate Connectors and Relay.

Audit-logs every routing-mode change on a Resource, making bypass decisions explicit and traceable rather than invisible gaps.

Userspace Networking NOTES ↗ STABLE

Twingate Linux Client adds userspace networking via HTTP Proxy Mode — no root or elevated privileges required.

New userspace networking mode for the Twingate Linux Client eliminates the need for root or elevated privileges.

> New HTTP Proxy Mode enables secure access to Twingate-protected resources over standard HTTP/HTTPS connections.
>
> Supports headless and restricted environments, enabling use in CI/CD pipelines, Docker containers, and developer workstations without admin access.

WAS THIS USEFUL? 👍 👎

Cohesive VNS3 Controller

Feedback ⓘ snapshot-20260713

NOTES ↗

VNS3 Controller is a virtual network appliance that secures and connects cloud infrastructure through encrypted overlays, firewalling, and network function virtualization.

VNS3 7.1.1 adds AI-assisted IPsec diagnostics and live plugin instance upgrades without disrupting running configurations.

└──▷ **HOW TO FIND IT**

Reach the AI Diagnostics advisor to troubleshoot a failing site-to-site IPsec tunnel without opening a support ticket.

> 📍 In the VNS3 Controller UI, navigate to the target IPsec connection, select the `Advice` option, enter your Cohesive auth token when prompted, and submit your IPsec configuration question to receive AI-generated diagnostics.

Introduces Connection Advisor with AI Diagnostics (Beta): an `Advice` option on VNS3 Controllers lets administrators ask for help configuring the controller or troubleshooting site-to-site IPsec connections (requires Cohesive auth token and active support contract).

Enables in-place plugin instance upgrades: a running plugin's full configuration can now be migrated to a new instance/version on the same VNS3 controller, starting at the same network address once the old instance is stopped.

WAS THIS USEFUL? 👍 👎

Cisco Duo Desktop

Feedback ⓘ snapshot-20260713 NOTES ↗

Cisco Duo Desktop is an endpoint security application that verifies device health and security posture before granting users access to protected resources.

Duo Desktop adds Microsoft Defender for Endpoint detection on Linux, Ubuntu 26.04 LTS support, and major macOS agent-detection performance gains.

Adds detection for Microsoft Defender for Endpoint on Linux (v4.6.0).

Supports Ubuntu 26.04 LTS on Linux; drops support for Ubuntu 23.04 and 20.04 LTS (v4.6.0).

Significantly improves performance of security agent detection on macOS (v7.18.0.0).

Prepares internal groundwork for upcoming Cynet and Check Point security agent detection on macOS (v7.18.0.0).

└─▷ **BREAKING ON UPGRADE**

! Support for Ubuntu 23.04 and Ubuntu 20.04 LTS is dropped in Duo Desktop for Linux version 4.6.0.

! Support for macOS 10.15 is removed in Duo Desktop for macOS version 7.5.0.0.

WAS THIS USEFUL? 👍 👎

DNSFilter Feedback ⓘ

06/30 June 30 – Web Application and API Update (2026-06-30)

5 RELEASES · 2026-06-15 → 2026-06-30 BLOG ↗ LATEST STABLE

DNSFilter is a cloud-based DNS security platform that blocks malicious websites, filters harmful content, and protects networks from cyber threats in real time.

DNSFilter Partner Portal relaunched with magic-link access and local lesson progress tracking.

▾ 4 more releases in this issue · 2026-06-15 → 2026-06-30

06/24 June 24 – Web Application and API Update (2026-06-24) BLOG ↗ STABLE

DNSFilter adds a CyberSight Agent compatibility warning banner to flag outdated Roaming Client versions.

New compatibility warning banner in CyberSight alerts when the Roaming Client version on a device is incompatible with the current CyberSight feature set, prompting an upgrade before relying on CyberSight data.

06/23 June 23 - macOS Roaming Client v2.4.2 Production (2026-06-23) BLOG ↗ STABLE

macOS Roaming Client v2.4.2 adds component version mismatch detection to catch post-upgrade inconsistencies.

└──▷ HOW TO FIND IT

After an MDM-managed upgrade, verify all components are in sync and resolve any mismatch without a reinstall.

📍 In the macOS menu bar, click the DNSFilter icon → if a version mismatch is detected, select the prompted option to resolve it directly from the menu.

Adds component version mismatch detection: alerts users when the menu bar app, system extension, and daemon are running mismatched versions — with an in-menu resolution option.

06/15 June 15 – Windows Roaming Client v3.4.23 Production (2026-06-15) BLOG ↗ STABLE

Windows Roaming Client v3.4.23 adds automatic Loopback mode fallback on Windows Server and more consistent device identification.

Adds automatic fallback to Loopback filtering mode when installed on Windows Server OS, preventing unexpected behavior on unsupported Server deployments.

Improves device identification consistency in the dashboard, reducing duplicate entries caused by missing identifiers after upgrades or reinstalls.

Enhances installer prerequisite validation and upgrade handling so failed upgrades no longer leave endpoints without the Roaming Client.

06/15 June 15 – macOS Roaming Client v2.4.2 Beta (2026-06-15) BLOG ↗

PRE-RELEASE

macOS Roaming Client v2.4.2 Beta adds component version mismatch detection to surface upgrade and MDM install issues.

└──▷ **HOW TO FIND IT**

After an MDM-managed upgrade, verify all Roaming Client components are in sync and resolve any mismatch without manual reinstall.

> 📍 In the menu bar app, open the DNSFilter menu — if a version mismatch is detected across the menu bar app, system extension, or daemon, an alert will appear with a prompt to resolve it directly from the menu.

Adds component version mismatch detection that alerts users when the menu bar app, system extension, and daemon are running different versions — with an in-menu resolution option.

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ macOS / iOS 1.5.18 (2026-06-24) BLOG ↗

Firezone Changelog tracks and publishes updates, new features, and bug fixes for Firezone, an open-source zero-trust network access and VPN management platform.

Firezone macOS/iOS 1.5.18 brings major UDP and TUN throughput upgrades, cutting CPU usage and packet loss on high-traffic connections.

Adds client-side evaluation of resource traffic filters, dropping disallowed packets locally to prevent repeated flow-creation request loops.

Batches TUN device reads and writes into a single syscall per batch on macOS and iOS, improving throughput and reducing CPU on high-traffic connections.

Introduces a per-destination UDP socket pool on macOS and iOS that unlocks Darwin's UDP fast path, enabling send batching and kernel congestion (flow advisory) signals.

Increases the macOS and iOS TUN read buffer size so the kernel can queue more outbound packets between reads, improving upload throughput.

Retries writes to the TUN device when its queue is momentarily full instead of dropping packets, reducing packet loss under bursty load on macOS and iOS.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

iDenfy

Feedback ⓘ snapshot-20260713 NOTES ↗

iDenfy is an identity verification platform that uses AI-powered document scanning and biometric facial recognition to automate KYC compliance for businesses.

iDenfy adds Brazil Digital CNH, new India/Philippines digital IDs, KYB automation API, tax ID fields, Risk Assessment PDF export, and more.

Supports Brazil Digital CNH (digital driver's license) as a verifiable document type.

Adds new digital ID support for India and the Philippines.

New API endpoint to re-run KYB automation workflows programmatically.

Adds dedicated tax identification fields — VAT Number, EIN, and TIN — within KYB workflows.

Introduces KYB Verify Later with automated IDV reminder notifications.

+7 more

WAS THIS USEFUL? 👍 👎

Passbolt

Feedback ⓘ Big Jet Plane | Browser Extension & API

8 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Passbolt Changelog documents version history, updates, bug fixes, and new features for Passbolt, an open-source password manager designed for team collaboration.

Passbolt 5.6.0 adds standalone notes, shared metadata key rotation, and resizable sidebars.

New standalone notes capability lets users create and store notes independently of password entries.

New shared metadata key rotation enables teams to rotate the encryption keys used for shared resource metadata.

New resizable sidebars let practitioners adjust the UI layout to surface more context during investigations or vault management.

▾ **7 more releases in this issue** · seen 2026-07-13

> Bloom | Browser Extension & API NOTES ↗ STABLE
>
> Passbolt 5.7.0 adds secret history, improved user/group management, and a detailed import summary report.
>
> New secret history feature lets practitioners review past versions of stored secrets.
>
> Improved user and group management capabilities in the browser extension and API.
>
> New detailed import summary report surfaces per-record results after bulk credential imports.

> Everything in its Right Place | Browser Extension & API NOTES ↗ STABLE
>
> Passbolt 5.8.0 adds dynamic RBAC role management, drag & drop group assignment, and maintenance updates.
>
> New dynamic role management in RBAC, enabling administrators to adjust role permissions without static configuration.
>
> New drag & drop interface for assigning users to groups, streamlining group membership management.

> Flight Facilities | Browser Extension & API NOTES ↗ STABLE
>
> Passbolt 5.9 adds PHP 8.5 support, account-discovery hardening, and DB end-of-life deadline visibility.

Adds PHP 8.5 support for teams preparing future infrastructure upgrades.

Strengthens protection against low-effort account discovery attacks.

Strengthens protection against UI spoofing/trickery attacks.

Surfaces upcoming database support deadlines more prominently in the platform.

Got To Be Real | Browser Extension & API NOTES ↗ STABLE

Passbolt 5.11 adds SMTP OAuth for Exchange Online, PingOne SSO, and production-ready SCIM provisioning.

└──▷ HOW TO FIND IT

Connect Passbolt mail delivery to Exchange Online using OAuth instead of a service-account password.

> 📍 In the Admin console, go to Administration › Email Server, select Microsoft Exchange Online as the provider, choose OAuth 2.0 as the authentication method, and enter your Azure app Client ID and Client Secret.

Enable PingOne as your SSO identity provider to let users authenticate via your PingOne directory.

> 📍 In the Admin console, go to Administration › Single Sign-On, select PingOne as the provider, and enter your PingOne Environment ID, Client ID, and Client Secret.

Enable production SCIM provisioning to automate user lifecycle management from your identity provider.

> 📍 In the Admin console, go to Administration › SCIM, enable the SCIM integration, and copy the SCIM base URL and bearer token into your IdP's provisioning settings.

Adds SMTP OAuth 2.0 support for Microsoft Exchange Online, enabling modern auth mail delivery without stored passwords.

Introduces PingOne SSO integration, expanding supported identity providers for single sign-on.

Promotes SCIM user provisioning out of beta, making it generally available for production directory sync.

I Feel Love | Browser Extension & API NOTES ↗ STABLE

Passbolt 5.13.0 adds in-app edition management, letting admins switch between Community and Pro editions.

└──▷ HOW TO FIND IT

Upgrade from Community to Pro edition without leaving the app.

📍 In the Passbolt admin console, go to Administration › Edition Management and select your target edition to switch.

Natural Blues | Browser Extension & API NOTES ↗ STABLE

Passbolt 5.12 launches Safari extension GA, adds PIN code resource type, and improves TOTP field detection.

Safari browser extension is now generally available for all users.

New dedicated PIN code resource type for storing numeric credentials separately from passwords.

Improved TOTP field detection for more reliable one-time password autofill.

You've Got the Love | Browser Extension & API NOTES ↗ STABLE

Passbolt 5.10 debuts the first Safari-compatible browser extension, bringing the password manager to Apple's browser.

└──▷ **HOW TO FIND IT**

Install the Passbolt extension on Safari to access and share passwords directly from Apple's browser.

📍 In Safari, go to the App Store, search for `Passbolt`, install the extension, then enable it under Safari › Settings › Extensions.

New Safari browser extension support: Passbolt is now compatible with Apple's Safari browser for the first time.

WAS THIS USEFUL? 👍 👎

RoboForm

Feedback ⓘ Version 9.9.4 on 2026-04-22 NOTES ↗

RoboForm is a password manager for Windows that securely stores, organizes, and autofills passwords and online forms across websites and applications.

RoboForm 9.9.4 adds Dashlane Safenotes import capability.

WAS THIS USEFUL? 👍 👎

Didit

Feedback ⓘ snapshot-20260713 NOTES ↗

Didit is a decentralized identity verification platform that provides KYC, authentication, and credential management tools for developers to integrate into their applications.

Didit June 2026: Document AI custom verification, MCP Server v5, PACE NFC on Android, new transaction rule operators, and brand import.

└──▷ **TRY IT**

Use the decision endpoint to retrieve AI-extracted fields from a custom document type after a verification session completes.

```
$ curl -X GET https://api.didit.me/v1/sessions/{session_id}/decision \
  -H "Authorization: Bearer <org_scoped_user_token>" \
  -H "Accept: application/json"
```

Install the Didit MCP Server v5 in Cursor so AI agents can run verifications and screen wallets using scoped access tokens.

📍 In Cursor, open Settings › MCP Servers › Add Server, search for `Didit` in the marketplace, and install the Didit plugin. Authenticate with your credentials — the server exchanges them for a narrowly-scoped token automatically.

Add a regex-based rule condition to flag transactions whose reference field matches a suspicious pattern, using the new regex operator in the Console rule editor.

📍 In the Business Console, go to Transaction Rules › New Rule › Add Condition, set the field to your target attribute, select the `regex` operator, and enter your pattern (e.g. `^[A-Z]{2}\d{6}$`). Save and publish the rule.

New Document AI Verification: define custom document types (contracts, licenses, certificates) with configurable extraction fields — no template required — and results are returned via decision/events endpoints and webhooks.

Auto-suggest extraction fields by uploading a sample document; the AI proposes fields for you directly in the workflow builder.

Document AI results surface in a dedicated Console section on session details, showing every extracted field, and full-name fields can be matched against the verified identity.

Document AI custom verification step available natively on iOS, Android, React Native, and Flutter (SDK 4.0.7+).

PACE NFC support on Android: passports and IDs protected with PACE can now be read over NFC.

+21 more

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ

Create Jira tickets from your Nudge Security findings

10 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Nudge Security is a SaaS security platform that discovers, manages, and governs all cloud and SaaS apps used across an organization's workforce.

Nudge Security adds native Jira integration to create, map, and auto-close tickets directly from security findings.

└──▷ HOW TO FIND IT

Connect Jira to Nudge Security so your team can route findings into their existing ticketing workflow.

📍 In the Nudge Security console, go to Settings › Ticketing and connect your Jira instance.

New Jira integration creates tickets directly from any finding, pre-filled with finding details and fix recommendations.

Supports mapping tickets to Jira project, work type, priority, assignee, and reporter during creation.

Maintains a linked audit trail between Nudge Security findings and their associated Jira issues.

▾ 9 more releases in this issue · seen 2026-07-13

Discover AI agents from the browser NOTES ↗ STABLE

Nudge Security can now discover AI agents built by employees on platforms without a public API, via its browser extension.

└──▷ HOW TO FIND IT

Enable browser-based agent discovery to start inventorying AI agents your employees are building across unsupported platforms.

📍 In the Nudge Security console, go to your browser extension settings and toggle on browser-based agent discovery. (Requires the Nudge Security browser extension to be deployed and enrollment in the AI agent discovery research preview.)

Request access to the AI agent discovery research preview if not yet enrolled.

📍 In the product navigation, go to Identities › AI Agents tab and select the option to request research preview access.

Discovers AI agents employees create across platforms lacking public APIs — including Cursor automations, Retool, Zapier Agents, ChatGPT workspace agents, OpenAI Agents Workflows, Atlassian Rovo (Jira, Bitbucket), Zoom AI Workflows, and HyperAgent — via the browser extension.

Surfaces discovered agents in the AI agents inventory under Identities, showing creator, components, and risk insights for each agent.

`Discover and revoke risky external file shares in Google Drive` NOTES ↗

STABLE

Nudge Security adds Google Drive external file share discovery and revocation in open beta.

└──▷ **HOW TO FIND IT**

Review all externally shared files and Shared Drives across your Google Workspace tenant to identify and revoke risky shares.

> 📍 In the Nudge Security console, go to Inventory › External File Shares. Connect the Google Drive connected app (one-time setup with in-product instructions), then use domain filters to prioritize cleanup by external company and take per-item revocation actions.

New Google Drive integration surfaces every externally shared file, folder, and Shared Drive across a Google Workspace tenant in a single view.

Enables filtering external shares by domain to see everything shared with a specific external company for targeted cleanup.

Per-item access controls let administrators tighten or revoke external share permissions with a confirmation step that details scope before changes are applied.

`Security posture findings now mapped to compliance frameworks` NOTES ↗

STABLE

Nudge Security maps security posture findings to SOC 2, NIST CSF 2.0, ISO 27001, HIPAA, and OWASP Top 10 Agentic AI Security frameworks.

└──▷ **HOW TO FIND IT**

Quickly scope your SOC 2 audit prep by filtering all posture findings to only those that violate SOC 2 controls.

> 📍 In the console, go to Security Posture › Findings, then use the compliance framework filter to select SOC 2. Review each listed finding and follow guided remediation from the detail panel.

Generate an audit-ready export covering all compliance frameworks for submission to an auditor or vCISO.

> 📍 In the console, go to Security Posture › Findings, apply any desired framework filters, then select Export to download a report with compliance framework and control columns included.

Maps security posture findings to five compliance frameworks: SOC 2, NIST CSF 2.0, ISO 27001, HIPAA, and OWASP Top 10 Agentic AI Security.

Adds compliance framework filtering on the Security Posture Findings view to surface every control violation across a SaaS estate at a glance.

Enables drill-down into any finding to view the specific controls it violates alongside guided remediation in the same panel.

Exports audit-ready reports with compliance framework and control columns included.

`Strengthen your SaaS security posture with our latest connected apps for Automox, Beyo`

NOTES ↗ STABLE

Nudge Security adds connected app integrations for Automox, BeyondTrust, and Ping Identity

New Automox connected app enables centralized visibility into accounts, activity, and misconfigurations for the endpoint management platform.

New BeyondTrust connected app brings privileged access management activity and account data into the Nudge Security SaaS inventory.

New Ping Identity connected app surfaces identity provider accounts and misconfigurations alongside the rest of your SaaS stack.

`Strengthen your SaaS security posture with our latest connected apps for Clockify, Acr`

NOTES ↗ STABLE

Nudge Security adds connected app integrations for Clockify, Acronis, and OneLogin.

Adds connected app integration for Clockify, enabling centralized visibility into accounts, activity, and misconfigurations.

Adds connected app integration for Acronis, enabling centralized visibility into accounts, activity, and misconfigurations.

Adds connected app integration for OneLogin, enabling centralized visibility into accounts, activity, and misconfigurations.

Strengthen your SaaS security posture with our latest connected apps for GitLab, Infor

NOTES ↗ STABLE

Nudge Security adds connected app integrations for GitLab, Informatica, and Logz.io

└──▷ **HOW TO FIND IT**

Enable the GitLab connected app in Nudge Security to surface account misconfigurations and activity across your GitLab organization.

> 📍 In the Nudge Security console, go to Integrations, search for GitLab, and follow the connection prompts to authorize and link your GitLab organization.

Adds connected app integration for GitLab, enabling centralized visibility into accounts, activity, and misconfigurations.

Adds connected app integration for Informatica, extending SaaS security posture coverage to the data integration platform.

Adds connected app integration for Logz.io, bringing log management and observability tooling into unified SaaS visibility.

Strengthen your SaaS security posture with our latest connected apps for Greenhouse, C

NOTES ↗ STABLE

Nudge Security adds connected apps for Greenhouse, Codacy, MuleSoft, and Oracle Cloud.

Adds connected app integration for Greenhouse to centralize visibility into accounts, activity, and misconfigurations.

Adds connected app integration for Codacy to centralize visibility into accounts, activity, and misconfigurations.

Adds connected app integration for MuleSoft to centralize visibility into accounts, activity, and misconfigurations.

Adds connected app integration for Oracle Cloud to centralize visibility into accounts, activity, and misconfigurations.

Strengthen your SaaS security posture with our latest connected apps for LastPass, Bam

NOTES ↗ STABLE

Nudge Security adds connected app integrations for LastPass, BambooHR, and Trello.

Adds connected app integration for LastPass, enabling centralized visibility into accounts, activity, and misconfigurations.

Adds connected app integration for BambooHR, enabling centralized visibility into accounts, activity, and misconfigurations.

Adds connected app integration for Trello, enabling centralized visibility into accounts, activity, and misconfigurations.

Strengthen your SaaS security posture with our latest connected apps for SentinelOne,

NOTES ↗ STABLE

Nudge Security adds connected app integrations for SentinelOne, Druva, GitGuardian, Kong, and Mural.

New SentinelOne connected app integration for centralized visibility into accounts, activity, and misconfigurations.

New Druva connected app integration for centralized visibility into accounts, activity, and misconfigurations.

New GitGuardian connected app integration for centralized visibility into accounts, activity, and misconfigurations.

New Kong connected app integration for centralized visibility into accounts, activity, and misconfigurations.

New Mural connected app integration for centralized visibility into accounts, activity, and misconfigurations.

WAS THIS USEFUL?  

ConductorOne

Feedback ⓘ July 10, 2026 (2026-07-10)

5 RELEASES · 2026-06-18 → 2026-07-10 BLOG ↗ LATEST STABLE

ConductorOne is an identity governance and access management platform that automates user access reviews, provisioning, and deprovisioning across enterprise applications.

ConductorOne adds real-time access provisioning, catalog indirect-access search, bulk revoke, visual condition builder, and Terraform connector ownership resources.

└──▷ **USE IT**

Migrate connector ownership in Terraform to avoid drift after the user_ids field is dropped from conductorone_integration_* resources.

hcl

```
resource "conductorone_connector_owner_user" "example" {
  connector_id = conductorone_integration_aws.example.connector_id
  user_id      = "<user-id>"
}

resource "conductorone_connector_owner_entitlement" "example" {
  connector_id   = conductorone_integration_aws.example.connector_id
  entitlement_id = "<entitlement-id>"
}
```

Connect an MCP server without managing client credentials when the server supports OAuth Dynamic Client Registration.

📍 In the C1 console, go to AI › MCP, add a new hosted or external MCP server, and enable Use dynamic client registration — C1 will handle OAuth registration automatically with no client ID or secret required.

Filter which Directory source accounts are imported into C1 using the visual condition builder with a live preview before saving.

📍 In the C1 console, go to Directory sources settings, open the three-dot menu next to your directory, select Condition, build your import filter visually, then click Preview to verify the match count and spot-check individual users before saving.

Adds real-time streaming of access profile grants and removals to connected apps, replacing hourly sync cycles.

New rolling-hour anomaly detector for membership removals, with a viewable and adjustable threshold in automation settings.

Request catalog now surfaces the role, team, or group that grants access to a searched resource, even when the entitlement name doesn't match the search term.

Enables bulk revoke of entitlements directly from the App Accounts page — select one or more accounts and remove specific entitlements or all access at once.

New visual condition builder for Directory source import filters replaces raw CEL expressions, with a built-in Data Preview to verify match counts before saving.

+15 more

└──▷ **BREAKING ON UPGRADE**

! The `user_ids` field on all `conductorone_integration_*` Terraform resources will no longer be returned by the Terraform provider starting in provider v1.4.0; if `user_ids` is configured in your Terraform files, this will cause drift and requires migration to the new `conductorone_connector_owner_user` and `conductorone_connector_owner_entitlement` resources and removal of the `user_ids` field.

! In the C1 navigation, User data sources in the Directory section is renamed to Directory sources and Groups is renamed to Dynamic groups; Inventory has moved from Explore to the new top-level Identity security section; MCP server configuration and C1 Gateway settings have moved under the new top-level `AI` section.

▾ **4 more releases in this issue** · 2026-06-18 → 2026-07-10

Private bridge for MCP servers (2026-06-26) BLOG ↗ STABLE

ConductorOne now governs private/air-gapped MCP servers via an outbound bridge agent — no inbound firewall rules required.

└──▷ **HOW TO FIND IT**

Deploy the bridge agent next to a private MCP server so its tools are governed by C1 without exposing inbound network access.

📍 In the C1 console, ensure AI access management is enabled for your tenant, then follow Settings › Connect a private MCP server through a bridge to deploy the bridge-client agent alongside your private MCP server.

Adds private MCP server connectivity via a lightweight `bridge-client` agent that maintains an outbound connection to C1, eliminating the need to open inbound firewall rules or expose servers to the public internet.

Enables full C1 entitlement governance (access requests, approvals, reviews, and auditing) for MCP tools running on-prem, in air-gapped clusters, or behind firewalls.

June 26, 2026 (2026-06-26) BLOG ↗ STABLE

ConductorOne adds single-page role mining, custom email providers, bulk role assignment, and flexible table exports.

└──▷ **HOW TO FIND IT**

Surface all denied or errored tasks in a large campaign without paging through the full task list.

> 📍 In the console, go to Campaigns › [your campaign] › Tasks, then open the Outcome filter and select Denied or Error to isolate tasks that need follow-up.

Onboard a new team by assigning the same C1 role to many users at once instead of editing each profile individually.

> 📍 In the console, go to Settings › Users, select the checkboxes for each new team member, and choose the bulk Change role action to assign their role in one step.

Produce a clean, auditor-ready policy export containing only the columns your compliance team needs.

> 📍 In the console, navigate to a policy's detail page or the policy list page and click Export to CSV to download the policy rules directly.

Redesigns role mining as a single-page experience with a token-based filter bar, inline cohort refinements, and a floating action bar for creating or updating access profiles without navigating away.

Adds early-access custom email provider support (Google Workspace, AWS SES, Microsoft 365, SendGrid) so C1 notifications send from your own domain instead of no-reply@conductorone.com.

Enables bulk role assignment — select multiple users on the Users page and apply a role change to all of them at once.

Introduces customizable, reorderable columns (saved per table) and targeted CSV export for the Apps, App accounts, Entitlements, and Resources tables.

Adds outcome-based filtering on campaign task lists (granted, denied, revoked, certified, error, canceled, timed out) to quickly surface tasks needing attention.

+7 more

Ask C1AI (2026-06-24) BLOG ↗ STABLE

ConductorOne launches Ask C1AI: a floating AI chat assistant for querying access data and taking action from anywhere in the platform.

└──▷ **HOW TO FIND IT**

While reviewing a policy page, open Ask C1AI to interrogate entitlements in context and act on findings without leaving the page.

> 📍 In C1, navigate to any page (e.g. an application or policy view) › click `Ask C1AI` in the bottom-right corner › use the context toggle to include the current page context › type your question about access or entitlements.

New Ask C1AI floating chat panel available on every page lets practitioners query org access data and take action without navigating away from their current view.

AI assistant ingests page context — the application, entitlements, or policy currently in view — to deliver more precise answers, with a toggle to include or exclude that context per conversation.

`Enterprise-managed authorization (2026-06-18)` **BLOG ↗** **STABLE**

ConductorOne adds enterprise-managed authorization for AI agents, issuing short-lived scoped tokens via the open XAA standard.

New early-access Enterprise-managed authorization enables AI agents (Claude, VS Code, etc.) to access organization-run MCP servers using short-lived, scoped tokens — no per-tool prompts, no persistent secrets.

Agents receive tokens only after entitlement checks, with every token request recorded for auditability.

AI agent access follows the same request, approval, review, and revocation workflows used for all other access in ConductorOne.

Built on the open Cross-App Access (XAA) standard for interoperability.

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Third-Party Apps for Organizations is now Generally Available (2026-07-10)

7 RELEASES · 2026-06-15 → 2026-07-10 BLOG ↗ LATEST STABLE

Auth0 is an identity and access management platform that provides authentication, authorization, and user management services for applications via APIs and SDKs.

Auth0 Organizations now supports third-party app access with per-org controls and scoped user consent.

.

└─▷ TRY IT

Promote a connection to domain level so a third-party app can authenticate users within a specific organization.

```
$ curl --request PATCH \
  --url
'https://<tenant>.auth0.com/api/v2/connections/<connection_id>' \
  --header 'Authorization: Bearer <mgmt_api_token>' \
  --header 'Content-Type: application/json' \
  --data '{"is_domain_connection": true}'
```

Enable third-party app access for a specific organization after opting in (access is blocked by default).

> 📍 In the Auth0 Dashboard, go to Organizations › <your-org> › Settings and enable third-party application access for that organization.

Enables third-party application access for Auth0 Organizations, configurable on a per-organization basis by tenant admins.

Blocks third-party app access by default for all existing organizations — admins must explicitly opt in, preserving backward compatibility.

Supports domain-level connections (`is_domain_connection: true`) as a requirement for authenticating org users through third-party apps.

Scopes user consent per organization — consent granted in one organization does not propagate to others.

▾ 6 more releases in this issue · 2026-06-15 → 2026-07-10

> Google One Tap Support for Universal Login EA (2026-07-09) BLOG ↗ STABLE

Auth0 Universal Login now supports Google One Tap for frictionless, one-click sign-in (Early Access).

└──▷ **HOW TO FIND IT**

Enable Google One Tap on your web app via the Auth0 Google Social Connection settings — navigate to the configuration documented in the Auth0 Google Social Connection docs.

> 📍 In the Auth0 Dashboard, go to Authentication › Social › Google › Settings and enable the Google One Tap option for your Universal Login prompt.

Adds Google One Tap support to the Universal Login prompt (Early Access), letting users authenticate instantly via a non-intrusive identity overlay without manual forms or page redirects.

`IPSIE session_expiry Claim Support for Okta and OIDC Enterprise Connections (2026-07-0`

BLOG ↗ **STABLE**

Auth0 adds IPSIE session_expiry claim support for Okta and OIDC Enterprise connections to enforce federated session ceilings.

└──▷ **USE IT**

Propagate the enforced session expiry to downstream apps so they can respect the IdP-driven session ceiling end-to-end.

javascript

```javascript
// Post-Login Action
exports.onExecutePostLogin = async (event, api) => {
  const sessionExpiry = event.session?.expiry;
  if (sessionExpiry) {
    api.idToken.setCustomClaim('session_expiry', sessionExpiry);
  }
};
```

Enable session_expiry enforcement on an existing OIDC Enterprise connection via the Management API.

```
$ curl -X PATCH https://<your-
tenant>.auth0.com/api/v2/connections/<connection_id> \
  -H 'Authorization: Bearer <mgmt_api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"options": {"session_expiry": true}}'
```

Enable session_expiry support on an Okta/OIDC Enterprise connection through the Auth0 Dashboard.

📍 In the Auth0 Dashboard, go to Authentication › Enterprise, select your Okta or OIDC connection, open its settings, and enable the Session Expiry option.

Supports the IPSIE SL1 `session_expiry` claim for Okta and OIDC Enterprise connections, using the IdP-provided expiry as a hard ceiling on Auth0 session lifetime.

Automatically terminates Auth0 sessions when a federated user's upstream IdP session expires, eliminating stale federated sessions and residual access.

Enforces session expiry as the minimum of three values: the IdP `session_expiry` claim, the tenant's absolute session expiration, and any expiry set via Actions.

Enables propagation of the final `session_expiry` value to downstream applications via a Post-Login Action that injects it as a custom claim in the Auth0-issued ID token.

Configurable per Okta/OIDC Enterprise connection through the Auth0 Dashboard or Management API.

`Automate Downstream Provisioning with Outbound SCIM for Users via Event Streams (2026-`

BLOG ↗ STABLE

Auth0 adds an Outbound SCIM Action template for Event Streams, enabling infrastructure-free user provisioning to any SCIM 2.0 app.

└──▷ **HOW TO FIND IT**

Enable automated downstream user provisioning to a SCIM 2.0 app without hosting any webhook infrastructure.

📍 In the Auth0 Dashboard, go to Event Streams › Create Stream, select Auth0 Actions as the destination, then apply the Outbound SCIM provisioning template. Edit the script to map Auth0 attributes to your downstream SCIM schema before activating the stream.

New Outbound SCIM Action template for Event Streams pushes `user.created`, `user.updated`, and `user.deleted` events to any SCIM 2.0-compliant downstream application without custom infrastructure.

Supports SCIM 2.0 REST operations (`POST`, `PUT`, `DELETE`) translated directly from Auth0 user lifecycle events.

Template-driven customization lets practitioners map Auth0 attributes to downstream SCIM schemas and enable opt-in UPSERT (create-on-update) behavior for gradual user trickle-in.

Built-in fault tolerance handles API timeouts and retries for transient failures natively within the Action script.

Refresh Token metadata is now Generally Available (2026-06-22) BLOG ↗

STABLE

Auth0 Refresh Token Metadata is now GA: attach up to 25 custom key-value pairs to refresh tokens via Actions and the Management API.

└──▷ **USE IT**

Tag refresh tokens with device and region context at login so downstream services can make access decisions without re-querying the IdP.

javascript

```
// Post-Login Action
exports.onExecutePostLogin = async (event, api) => {
  api.refreshToken.setMetadata('deviceName',
event.request.user_agent);
  api.refreshToken.setMetadata('loginRegion',
event.request.geoip?.countryCode);
  api.refreshToken.setMetadata('orgContext',
event.organization?.id);
};
```

Inspect a refresh token's stored metadata via the Management API to audit contextual data tied to a specific token.

```
$ curl -X GET "https://<your-domain>/api/v2/refresh-tokens/<id>" \
  -H "Authorization: Bearer <mgmt_api_token>"
```

Update metadata on an existing refresh token — for example, to mark a token as elevated after step-up authentication.

```
$ curl -X PATCH "https://<your-domain>/api/v2/refresh-tokens/<id>" \
  -H "Authorization: Bearer <mgmt_api_token>" \
  -H "Content-Type: application/json" \
  -d '{"metadata": {"stepUpCompleted": "true"}}'
```

Adds Refresh Token Metadata (GA for Enterprise): attach up to 25 custom key-value pairs to refresh tokens that persist throughout the token lifecycle.

New api.refreshToken.setMetadata() method in Post-Login Actions lets you store contextual data (device, region, org) at login time.

New Management API endpoints support retrieving (`GET`), updating (`PATCH`), and revoking (`DELETE`) refresh tokens with metadata at `/api/v2/refresh-tokens/{id}` .

Google Workspace Directory Sync for Groups - Early Access Updates (2026-06-17)

BLOG ↗ STABLE

Google Workspace Directory Sync for Groups enters open Early Access with native RBAC role assignment for synced groups.

└──▷ **HOW TO FIND IT**

Assign a synced Google Workspace group to an Auth0 RBAC role so members automatically receive the correct permissions when they log in.

> 📍 In the Auth0 Dashboard, go to Authentication › Enterprise › Google Workspace connection › Synchronization, configure your group synchronization preferences, then navigate to User Management › Roles › <your role> › Groups and assign the synced Google Workspace group.

Enables Google Workspace Directory Sync for Groups in Early Access for all public cloud customers — no enrolment required.

Supports automated group synchronization, `Sync all`, and partial group selection out-of-the-box.

Adds native assignment of synced Google Workspace groups to Auth0 tenant-level RBAC roles and organization-specific roles, so group members automatically inherit permissions at login.

Dashboard Search for APIs Now in Beta (2026-06-15)

BLOG ↗ STABLE

Auth0 Dashboard now supports real-time API search by ID, identifier, or name — no more paginated scrolling.

└──▷ **HOW TO FIND IT**

Quickly locate a specific API in large tenants without paging through lists.

> 📍 In the Auth0 Dashboard, go to Applications › APIs and type the API name, identifier, or ID into the search bar to filter results in real time.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

Enhancement: SailPoint MCP Server Now Supports Global URL and OAuth — Connect Any AI Cli

15 RELEASES · 2026-06-16 → 2026-07-09 BLOG ↗ LATEST STABLE

SailPoint is an identity security platform that manages and governs user access, permissions, and authentication across enterprise applications and systems.

SailPoint MCP Server adds a global tenant-agnostic URL and OAuth 2.1/PKCE auth, plus native Cursor IDE support.

└──▷ **USE IT**

Connect Cursor IDE to SailPoint ISC so an AI agent can handle access requests without manual token management.

json

```
// ~/.cursor/mcp.json
{
  "mcpServers": {
    "sailpoint-access-requests": {
      "url": "https://mcp.api.cloud.sailpoint.com/latest/access-
requests/mcp",
      "auth": {
        "CLIENT_ID": "your-client-id",
        "scopes": ["sp:scopes:all"]
      }
    }
  }
}
```

Create the required public OAuth API client in ISC before wiring up Cursor or Claude Desktop.

📍 In the ISC console, go to Admin › Security Settings › API Management › Create API Client: - Set client type to Public - Enable Authorization Code + Refresh Token grants - Set Redirect URL to: cursor:/`/anysphere`.cursor-mcp/oauth/callback - Set Scope to: sp:scopes:all Copy the resulting Client ID for use in mcp.json.

New global MCP endpoint (`https://mcp.api.cloud.sailpoint.com/latest/access-requests/mcp`) works for all non-FedRAMP tenants — no per-org URL construction needed.

Replaces manual PAT/token workflow with OAuth 2.1 + PKCE: browser-based login, SSO support, consent screen, and automatic token refresh for up to one month.

Adds native Cursor IDE support: configure the global URL and Client ID in `~/.cursor/mcp.json`, and the full OAuth flow is handled automatically on connect.

AI assistants can now manage ISC access requests conversationally — search the access catalog, submit requests, check approval status, and cancel requests.

▾ **14 more releases in this issue** · 2026-06-16 → 2026-07-09

New Capability: Data Access Security (DAS) for DataBricks (2026-07-09) BLOG ↗

STABLE

SailPoint DAS adds a Databricks connector for resource discovery, permission analysis, and data classification.

New Databricks connector for Data Access Security (DAS) enables discovery of Databricks resources including Tables, Views, and Stored Procedures.

Supports permission analysis to detect over-permissioned accounts and limit entitlement sprawl across Databricks deployments.

Ingests data classification labels directly from Databricks (when enabled), surfacing classification forensics within the DAS platform (preview).

Unifies Databricks governance with existing SailPoint frameworks covering applications, infrastructure, and file storage systems.

Expanding PAM application discovery: CyberArk Enhancement and New BeyondTrust & Deline

BLOG ↗ STABLE

SailPoint adds BeyondTrust & Delinea PAM discovery connectors and upgrades CyberArk to instance-level application discovery.

└──▷ **HOW TO FIND IT**

Set up the new BeyondTrust Password Safe Cloud PAM Discovery connector to maintain a live inventory of managed systems in ISC.

> 📍 In the ISC console, go to Admin › Connections › Discovery Connectors › Create Connector › PAM category › Configure next to BeyondTrust Password Safe Cloud PAM. In Express Setup, enter your OAuth 2.0 client_id and client_secret from the Application user assigned to the group with Password Safe System Management (Read) permission, then click Test Connection and Discover (or set a recurring schedule under Additional Settings › Discovery Settings).

Set up the new Delinea Secret Server PAM Discovery connector to catalogue secrets and their metadata in ISC.

📍 In the ISC console, go to Admin › Connections › Discovery Connectors › Create Connector › PAM category › Configure next to `Delinea Secret Server PAM`. In Express Setup, paste the Secret Server URL from Settings › Secret Server › Secret Server Connection as Host URL and enter the Service User username/password as Client ID/Client Secret, then click Test Connection and Discover.

Validate the upgraded CyberArk connector returns instance-level results after granting Safe-level permissions to avoid a blank discovery run.

📍 In Privilege Cloud, go to Policies › Safes, open each relevant Safe › Members › Add Member, add the connector service user with Read-only + Access enabled permissions. Then in ISC, open the CyberArk Privilege Cloud Shared Services PAM connector and click Test Connection followed by a manual discovery aggregation before the next scheduled run.

Upgrades CyberArk Privilege Cloud Shared Services PAM Discovery connector to discover individual application instances (e.g., three MySQL instances as three distinct entries) rather than generic platform templates.

New BeyondTrust Password Safe Cloud PAM Discovery connector retrieves managed systems via the BeyondTrust public API using OAuth 2.0 client credentials, with support for continuous scheduled discovery.

New Delinea Secret Server PAM Discovery connector discovers secrets and associated metadata (names, IDs, template types, creation dates, active status) from Delinea Secret Server via OAuth 2.0, mapping results into ISC-consumable datasets.

Breaking change for existing CyberArk Discovery customers: next discovery aggregation after the update will return instance-level applications instead of platform templates; new Safe-level read permissions must be configured on the CyberArk service user before the next scheduled run or discovery may return no results.

└──▷ **BREAKING ON UPGRADE**

- ! Existing CyberArk PAM Discovery connector users will see a different set of discovered applications after the update: instance-level applications are now discovered instead of platform-level templates, so previously discovered platform-level entries may no longer appear and new instance-level entries will be returned instead.
- ! If the required CyberArk service user permissions are not configured before the next discovery run (including the Privilege Cloud Users role and Read-only Safe membership for each relevant Safe), the CyberArk PAM Discovery connector may return no applications.

`Enhancement: ServiceNow Service Catalog for Identity Security Cloud v5.3 Released (202`

BLOG ↗ STABLE

ServiceNow Service Catalog for ISC v5.3 adds Additional Owner approvals, questionnaire UX, and faster OAuth.

Supports `Additional Owner` approval type, expanding flexibility in organizational approval flow definitions.

Introduces `Additional Info / Questionnaire` capabilities for catalog item requests, streamlining the requester experience.

Reduces total OAuth calls to the Identity Security Cloud API, improving integration response times.

`New Capability: Privilege On Demand (2026-07-08)` BLOG ↗ STABLE

SailPoint adds Privilege On Demand (JIT-P v1.0) to shift elevated access from standing to time-bounded, on-demand activation.

└──▷ **HOW TO FIND IT**

Configure JIT activation windows and mark entitlements as eligible so users can activate privileged access on demand rather than holding it standing.

> 📍 In the Admin console, go to Admin Settings › Privilege on Demand and set global JIT activation duration and extension duration. Then navigate to the target entitlement, enable `Requestable`, and configure JIT entitlement assignment with the desired approval behavior (auto-approve or approval-required).

Introduces Privilege On Demand (JIT-P v1.0): an activation-based model where eligible entitlements are requested, then explicitly activated for a bounded duration with automatic expiration — replacing always-on standing access.

Enables configurable global JIT activation and extension durations via Admin settings, giving IAM teams control over how long elevated access windows last.

Supports auto-approve or approval-required workflows for JIT entitlement activation, selectable per entitlement.

Delivers end-user JIT activation through Launchpad, allowing authorized users to self-activate approved access when work begins.

Adds five new email notification templates: Access Request Assignment, JIT Activation 15-Minute Reminder, Just-In-Time Activation Extended, Just-In-Time Activation Failed, Just-In-Time Activation Ready, and Just-In-Time Deactivated — providing full lifecycle visibility for activation events.

`+1 more`

`New Capability: View Permission Paths Now Available for Campaign Reviewers (2026-07-02`

BLOG ↗ STABLE

SailPoint DAS adds in-context Permission Path viewer for campaign reviewers, eliminating navigation away from the review screen.

.
.
.

└──▷ **HOW TO FIND IT**

Review how a user obtained access to a sensitive resource without leaving the certification screen during a Permissions-type campaign.

> 📍 In the DAS console, go to My Tasks › Access Certification, select the desired Permissions-type campaign, locate the user row in the review table, and click the Permission Path icon in the Actions column to open the Permission Path overlay.

New `View Permission Path` action in the Actions column of the Campaign Review screen opens a pop-up showing the full permission path for a user's access, including whether access is actively used.

Supports visualization of multiple permission paths to a single resource within Permissions-type campaigns.

Permission Path icon is available for User permission rows and visible-but-disabled for Group permission rows, giving reviewers clear context on scope.

`Enhancement: Data Access Security Resource Selector Redesign (2026-07-02)`

BLOG ↗ **STABLE**

Data Access Security gets a redesigned Resource Selector with full-view modal and free-text search.

.

└──▷ **HOW TO FIND IT**

Use the new overlay to quickly locate a deeply nested cloud resource when building a Forensics permission filter.

> 📍 In the console, go to Forensics › Permissions, set Attribute to `Resource` and Operator to `Any Of`, then click the Value field — the new Browse and Select Resources overlay opens. Type a name or path in the search bar to jump directly to the target resource without expanding the folder tree manually.

Introduces a new `Browse and Select Resources` overlay in Data Access Security, replacing the previous compact resource picker with a full-view modal.

Adds free-text search bar (`Search by name or path`) to the Resource Selector, enabling instant navigation to resources without manually drilling through folder hierarchies.

New hierarchical tree view supports expanding and collapsing major integrations

(AWS, Azure, Box, GDrive, OneDrive) and nested drive structures within the selector.

Maintenance: HTTP/2 Enablement in Cloudflare (2026-06-30) BLOG ↗ STABLE

SailPoint ISC enables HTTP/2 via Cloudflare on Jul 27, 2026 — API response headers will shift to lowercase.

Enables HTTP/2 on all ISC environments via Cloudflare, delivering improved security and performance over HTTP/1.1.

Breaking change: HTTP/2 lowercases all response header keys (e.g., `X-Total-Count` becomes `x-total-count`) — scripts and API clients that match headers case-sensitively must be updated.

└──▷ **BREAKING ON UPGRADE**

! HTTP/2 enablement in Cloudflare will cause response header keys to become lowercase (e.g., `X-Total-Count` becomes `x-total-count`), which may break API calls or scripts that expect case-sensitive header keys in responses from ISC

New Capability: Customizable DAS Insights on ISC Homepage (2026-06-29) BLOG ↗

STABLE

SailPoint ISC homepage gains Data Access Security widgets and a pre-built DAS dashboard template for unified identity + data risk visibility.

└──▷ **HOW TO FIND IT**

Quickly stand up a unified DAS monitoring view for your security team without manually assembling individual widgets.

📍 In the ISC console, go to MySailPoint Homepage › Dashboards › Dashboard Templates and select `Data Access Security Dashboard` to instantiate the pre-built template.

Adds a suite of Data Access Security (DAS) widgets to the MySailPoint homepage, enabling instant visibility into unstructured data access exposure and high-risk accounts with excessive permissions.

Introduces a pre-built `Data Access Security Dashboard` template that bundles all available DAS widgets for quick, one-click setup.

Enables drill-down from any DAS widget directly into Data Access Security forensics for deeper analysis of flagged metrics.

Enhancement: Tablet Support for Certifications (2026-06-26) BLOG ↗ STABLE

SailPoint ISC adds dedicated tablet support for Certification Reviewers with smart device detection and landscape-mode prompting.

Adds automatic tablet detection for Certification Reviewers, evaluating screen capabilities to determine if the reviewer UX can be displayed optimally.

Introduces landscape orientation prompting on tablets, guiding reviewers to rotate their devices for full visibility of access review tables and critical details.

Delivers an optimized Certification Reviewer UX in landscape mode, presenting a comprehensive view of access items to support informed decision-making on mobile devices.

`Announcement: Salesforce External Client App is now Available (2026-06-24)`

`BLOG ↗` `STABLE`

SailPoint ISC Salesforce connector now supports External Client App (ECA) auth; legacy OAuth flows retire Sep 30, 2026.

└──▷ **HOW TO FIND IT**

Enable ECA authentication on an existing ISC Salesforce connector to maintain connectivity before the September 30, 2026 deadline.

> 📍 In the Identity Security Cloud UI, open the Salesforce connector › go to the ECA Authentication section › turn on `Enable External Client App (ECA)` › under Client Configuration, add key `host_name` with value `login.salesforce.com` (production) or `test.salesforce.com` (sandbox) › click Authorize and complete the OAuth 2.0 Authorization Code Grant flow › click Save › click Test Connection to verify.

Adds support for Salesforce External Client App (ECA) authentication as the new default and only supported auth method for the ISC Salesforce connector (VA and SaaS).

Supports OAuth 2.0 Authorization Code Grant flow exclusively, replacing JWT Bearer, Refresh Token, Client Credentials, and Basic Authentication.

New `Enable External Client App (ECA)` toggle and ECA Authentication section available in the Identity Security Cloud UI Salesforce connector configuration.

Breaking change: Basic Auth and all OAuth 2.0 Connected App flows are permanently removed October 1, 2026 — connectors not migrated will lose all Salesforce connectivity.

Backward compatibility toggle available in the connector UI during the 90-day migration window (July 1–September 30, 2026) to temporarily revert to legacy auth if issues arise.

└──▷ **BREAKING ON UPGRADE**

! Basic Authentication, OAuth 2.0 JWT Bearer, OAuth 2.0 Refresh Token, and OAuth 2.0 Client Credentials are permanently removed from the ISC Salesforce connector (VA and SaaS) on October 1, 2026; connectors not migrated to External Client App (ECA) Authentication will lose Salesforce connectivity and all connector operations (provisioning, aggregation) will fail.

`New Capability: Accelerated Data Discovery with Sampling Ratio (2026-06-24)`

BLOG ↗ STABLE

SailPoint Data Discovery adds configurable Sampling Ratio for faster, lighter sensitive-data classification scans.

.
.
.

└──▷ **HOW TO FIND IT**

Quickly surface sensitive-data hot spots in a newly integrated application before committing to a full scan.

📍 In the console, go to Admin › Applications › Add New, complete the application setup, then on the Data Classification configuration screen set the Sampling Ratio to your desired percentage (default is 10%).

Tune the sampling percentage on an existing application to balance discovery speed against scan coverage.

📍 In the console, go to Compliance › Data Classification › Application Scope, locate the application, click Edit, and adjust the Sampling Ratio sliding scale to the desired value.

New Sampling Ratio option for Data Classification scans lets you scan a configurable percentage of documents (default 10% for new apps) to rapidly identify sensitive-data hot spots without a full resource-intensive scan.

New applications are automatically onboarded with 10% sampling enabled by default, accelerating time-to-insight for sensitive data exposure.

Existing applications with Data Classification enabled retain 100% sampling by default to preserve consistency with prior behavior.

New Sampling Ratio sliding-scale control added to the Application Scope edit screen under Compliance > Data Classification > Application Scope.

`Enhancement: CIEM – Support for Cloud Resource Tags/Labels (2026-06-23)`

BLOG ↗ STABLE

SailPoint CIEM now surfaces AWS Tags, Azure Tags, and GCP Labels alongside cloud resource access data.

└──▷ **HOW TO FIND IT**

When you need to answer 'who has access to production-tagged cloud resources', use ISC Search to filter by tag key and value without leaving SailPoint CIEM.

> 📍 In ISC CIEM, go to Search and query by cloud resource tag key (e.g., Environment) and tag value (e.g., Production) to list all identities with effective access to matching resources.

Adds native AWS Tags, Azure Tags, and GCP Labels visibility directly within the CIEM effective access view, enabling practitioners to see which identities can access tag-scoped resources.

Enables ISC Search to query cloud resources by tag/label key and value, supporting ad-hoc access reviews such as 'who can access PII-tagged resources'.

Surfaces cloud resource tags and labels within CIEM AIC dashboards for reporting and governance validation across cloud environments.

Enhancement: Identity Graph UI Updates (2026-06-19) BLOG ↗ STABLE

Identity Graph UI redesigned with clearer visual hierarchy, elevated Data Table, and consolidated controls rail.

Colocates help menus, documentation, and the legend into a unified location for faster access.

Makes the graph's root node more visually apparent, improving at-a-glance orientation.

Elevates the Data Table to a first-class UI citizen alongside the graph view.

Redesigns Explorer to function as a search history rather than a navigation panel.

Consolidates filtering, grouping, and other widgets into a single horizontal rail.

+1 more

Deprecation: Virtual Appliance network access - AWS URLs (2026-06-16) BLOG ↗ STABLE

Virtual Appliance traffic migrating to SailPoint-hosted URLs for SQS and ECR; new hostnames to allowlist by June 29.

└──▷ **TRY IT**

Add the new SailPoint-hosted Docker/ECR endpoint to your firewall allowlist before the June 29 production rollout to ensure uninterrupted VA image updates.

```
$ # Commercial environments — add to your outbound firewall
allowlist:
va-docker.secure-api.infra.identitynow.com

# FedRAMP environments — add to your outbound firewall allowlist:
va-docker.secure-api.saas.sailpointfedramp.com
```

Introduces SailPoint-hosted endpoints for Virtual Appliance message queue traffic, replacing direct regional Amazon SQS URLs (`sqs.<region>.amazonaws.com`).

Introduces SailPoint-hosted endpoints for VA container image updates, replacing direct Amazon ECR paths — new hostnames: `va-docker.secure-api.infra.identitynow.com` (Commercial) and `va-docker.secure-api.saas.sailpointfedramp.com` (FedRAMP).

Sandbox rollout of SailPoint-hosted URLs begins June 15; production rollout begins June 29, 2026.

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ 08 Jul 2026 (2026-07-09) BLOG ↗

Ping Identity is an identity security platform that provides authentication, single sign-on, multi-factor authentication, and access management solutions for enterprises.

Ping Identity AIC adds social IdP management, self-service journey mapping, PKCS5S2 migration, and expanded ESV support in the admin console.

└──▷ **HOW TO FIND IT**

Manage Proxy Connect configuration through the UI instead of making API calls — useful for tenants who want to avoid direct API access for routine config changes.

> 📍 In the admin console, go to Settings › Proxy Connect and update your Proxy Connect configuration directly in the panel.

Use an ESV to centrally manage and rotate the `from` address for email providers without hardcoding values in the console.

> 📍 In the admin console, go to Email Provider settings, and set the `From Address` field to an ESV reference (e.g., & `esv.email.from.address`).

Create or edit a social identity provider entirely within the admin console, avoiding manual API calls.

> 📍 In the admin console, go to Identities › Social Identity Providers, then select Create Provider or edit an existing one.

Extends the Evaluator Version selector in the admin console to all remaining next-generation script types: OIDC node, Social Provider Handler node, OAuth2 scripts, and Policy Condition scripts.

Enables migration of user accounts from the deprecated PKCS5S2 password hashing scheme to a more secure algorithm.

Adds an Account Status column (Active/Inactive) to the Identities page.

Supports using an environment secret or variable (ESV) to configure the `from` address for email providers.

Allows Proxy Connect–enabled tenants to manage Proxy Connect configuration directly in the admin console without using the API.

+7 more

└──▷ **BREAKING ON UPGRADE**

! Microsoft 365 SSO applications now use the selected UPN Attribute Name setting for WS-Trust username authentication; existing applications must be resaved for this

change to take effect.

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password CLI Beta 2.37.0-beta.01 (2026-07-09)

16 RELEASES · 2026-06-18 → 2026-07-09 BLOG ↗ PRE-RELEASE

1Password is a password manager that securely stores, organizes, and autofills passwords, credentials, and sensitive information across devices and browsers.

1Password CLI Beta 2.37.0 adds 14 new shell plugins, tab completion for bash/zsh, and GitHub Enterprise Server support.

Adds Shell Plugin support for 14 new CLIs: Cline, CrateDB, Cursor, Descope, Exercism, Expo/EAS, GitHub Copilot, Google Gemini AI API, JetBrains Junie, Kiro, OpenCode, OpenTofu, PyPI (twine/flit/hatch), and UpCloud — enabling Touch ID and other unlock options for each.

Extends the GitHub shell plugin to provision `GH_ENTERPRISE_TOKEN` for GitHub Enterprise Server hosts and support simultaneous authentication to both github.com and Enterprise with separate credentials.

Adds environment variable-based provisioning support to the Redis CLI shell plugin.

Adds tab completion when using shell plugins in bash and zsh shells.

Enables Terraform workspace subcommands to trigger authentication when using 1Password Shell Plugins.

▾ **15 more releases in this issue** · 2026-06-18 → 2026-07-09

> 1Password for Android beta 8.12.28-17 (2026-07-01) BLOG ↗ STABLE
>
> 1Password for Android adds exact-phrase search with support for special characters like IPs and hostnames.
>
> Supports quoted exact-phrase search queries, enabling lookups of strings containing dots, dashes, and other special characters (e.g., `10.1.2.3`).

> 1Password for Linux beta 8.12.28-17 (2026-07-01) BLOG ↗ STABLE
>
> 1Password Linux beta adds exact-phrase search, simpler MCP server launch, and a unified Developer settings section.
>
> └──▷ **HOW TO FIND IT**
>
> Find a credential by exact IP address or hyphenated hostname without wildcard noise.
>
> > 📍 Search for `10.1.2.3` or `prod-db-east` in the 1Password search bar to retrieve exact matches.
>
> Reference the MCP server in a client config without hard-coding its install path.

```json
{
  "mcpServers": {
    "1password": {
      "command": "1password-mcp"
    }
  }
}
```

Access all developer integrations (CLI, SDKs, MCP) from one place.

> 📍 In the app, go to Settings › Developer to find 1Password CLI, SDKs, and MCP client integration grouped together.

Supports exact-phrase search queries wrapped in double quotes, enabling precise lookup of IP addresses, hostnames, and strings containing dots or dashes.

Enables launching the bundled `1password-mcp` binary by name instead of requiring an absolute path when configuring MCP server clients.

Consolidates 1Password CLI, SDKs, and MCP client integration into a single unified section under Settings > Developer.

`1Password for Mac beta 8.12.28-17 (2026-07-01)` BLOG ↗ STABLE

1Password beta adds exact-phrase search, native HTTP auth autofill, and a launchable MCP binary alongside developer settings consolidation.

└──▷ HOW TO FIND IT

Find a vault item by exact IP address or hyphenated string without partial-match noise.

> 📍 In the 1Password search bar, type `10.1.2.3` (with double quotes) to retrieve only items matching that exact string.

Configure an MCP client to use the bundled binary without hardcoding an install path.

```json
{
  "mcpServers": {
    "1password": {
      "command": "1password-mcp"
    }
  }
}
```

Supports exact-phrase search queries using double quotes, including special characters like IP addresses (e.g., `10.1.2.3`).

Adds Universal Autofill support for native HTTP authentication prompts in browsers.

Enables launching the bundled `1password-mcp` binary by name instead of requiring an absolute path for MCP server clients.

Consolidates 1Password CLI, SDKs, and MCP client integration into a single section under Settings > Developer.

`1Password for Windows beta 8.12.28-17 (2026-07-01)` BLOG ↗ STABLE

1Password for Windows adds exact-phrase search with special characters and a simpler MCP server launch path.

└──▷ **HOW TO FIND IT**

Find an item containing a specific IP address or dotted string without partial matches cluttering results.

> 📍 In the 1Password search bar, type `10.1.2.3` (with double quotes) to retrieve only items matching that exact string.

Configure an MCP client to use the bundled binary without hardcoding an install path — useful in shared or portable setups.

> 📍 In your MCP client config, set the server command to `1password-mcp` instead of an absolute path such as `C:\Users\you\AppData\Local\1Password\app\8\1password-mcp.exe`.

Adds exact-phrase search: wrap a query in double quotes to match special characters like dots and dashes (e.g., `10.1.2.3`).

Enables launching the bundled `1password-mcp` binary by name instead of requiring an absolute path in MCP client configs.

Groups 1Password CLI, SDKs, and MCP client integrations into a single unified section under Settings > Developer.

`1Password for iOS beta 8.12.28-17 (2026-07-01)` BLOG ↗ STABLE

1Password iOS beta adds exact-phrase search with special characters and password generation without unlocking via Autofill.

└──▷ **HOW TO FIND IT**

Look up a vault item by exact IP address or dotted identifier without partial-match noise — useful when managing infrastructure credentials.

> 📍 In the Items tab search bar, type `10.1.2.3` (with double quotes) to retrieve only items matching that exact string.

Generate a strong password from an app's login screen without leaving the Autofill flow or unlocking 1Password — reduces friction during account creation on-device.

> 📍 Tap a password field in any app › tap the keyboard accessory or Autofill prompt › select `Generate Password` — no 1Password unlock required.

Supports exact-phrase search using double quotes, including special characters like dots and dashes (e.g., `10.1.2.3`).

Enables password generation through Autofill without requiring 1Password to be unlocked first.

Improves Credential Exchange export error messages with actionable remediation steps.

`1Password in the browser 8.12.28-17 (2026-07-01)` **BLOG ↗** STABLE

1Password browser extension adds exact-phrase search and smarter phishing-prevention prompts.

Supports exact-phrase search queries using double quotes, including special characters like IP addresses and domain names (e.g., `10.1.2.3`).

On verified low-risk sites, phishing prevention prompt now promotes `Add website and autofill` as the primary action.

`1Password for Android 8.12.26 (2026-06-30)` **BLOG ↗** STABLE

1Password Android adds `Open 1Password` shortcut in autofill suggestions and updates item-move/duplicate audit events to Export Vault.

Adds an `Open 1Password` option directly in the autofill login suggestion list, letting users search for any vault item without leaving the sign-in flow.

Changes the audit event type for items moved or duplicated between 1Password accounts from a generic event to an Export Vault event, improving audit log fidelity for security teams.

`1Password for Linux 8.12.26 (2026-06-30)` **BLOG ↗** STABLE

1Password for Linux 8.12.26 adds configurable default landing section and updates audit events for cross-account item moves.

└──▷ HOW TO FIND IT

Set the app to open directly to a specific vault or section instead of the default Home view, reducing clicks for security teams managing multiple vaults.

> 📍 In the app, go to Settings › General › Defaults and select the desired section to show on launch.

Adds a configurable default landing section so the app opens to a chosen view after launch or inactivity (Settings > General > Defaults).

`1Password for Mac 8.12.26 (2026-06-30)` **BLOG ↗** STABLE

1Password 8.12.26 adds configurable default view, CTAP2 authenticator support, and improved Safari inline menu navigation.

└──▷ **HOW TO FIND IT**

Set the default landing section so your team opens directly to a specific vault or view each session.

> 📍 In the app, go to Settings › General › Defaults and select your preferred starting section.

Adds configurable default section shown on app open or resume (Settings > General > Defaults).

Supports CTAP2 authenticators requiring user verification for 1Password account sign-in with two-factor authentication.

Changes the audit event logged for items moved or duplicated between accounts to an Export Vault event.

Enables Safari inline menu keyboard navigation: pressing the down arrow in a field now focuses the first suggestion.

`1Password for Windows 8.12.26 (2026-06-30)` **BLOG ↗** STABLE

1Password for Windows 8.12.26 adds configurable default view, improved audit events, and smarter display handling.

└──▷ **HOW TO FIND IT**

Set the landing view so your team's shared vault is front-and-center every time the app opens.

> 📍 In the app, go to Settings › General › Defaults and choose your preferred section to display on startup.

Adds configurable default section shown on app open or resume via Settings > General > Defaults.

Changes the audit event for items moved or duplicated between accounts to an Export Vault event, giving security teams more accurate audit trails.

Large Type and QR code windows now open on the active display rather than the primary display.

`1Password in the browser 8.12.26 (2026-06-30)` BLOG ↗ STABLE

1Password browser extension 8.12.26 adds instant unlock fallback, system language matching, and body-level field exclusion attributes.

Enables immediate unlock of the browser extension when the 1Password desktop app is not detected, eliminating the previous forced connection-attempt delay.

Adds `Use system defaults` option in the language selector so the extension matches the browser's configured language automatically.

Supports body-level `data-1p-ignore` and `data-op-ignore` attributes to exclude entire pages' fields from 1Password autofill.

Improves policy-change detection for enterprise/MDM-managed deployments affecting the browser extension.

`1Password for Android beta 8.12.26-35 (2026-06-23)` BLOG ↗ STABLE

Android beta adds `Open 1Password` shortcut in autofill suggestions and changes item-move audit events to Export Vault.

Adds an `Open 1Password` option directly within the autofill suggested-logins list, letting users search for any item without leaving the sign-in flow.

Changes the audit event recorded when an item is moved or duplicated between 1Password accounts from a generic event to an Export Vault event, improving audit trail clarity for security teams.

`1Password for Android beta 8.12.26-31 (2026-06-18)` BLOG ↗ STABLE

1Password Android beta adds quick vault access in autofill and updates audit event for item moves/duplicates.

Adds an `Open 1Password` option in the autofill suggested-logins list, letting users search the full vault without leaving the sign-in flow.

Changes the audit event logged when an item is moved or duplicated between 1Password accounts from a generic event to an Export Vault event.

`1Password for Linux beta 8.12.26-31 (2026-06-18)` BLOG ↗ STABLE

1Password Linux beta adds configurable default section on launch and updates audit event classification for item moves.

└──▷ **HOW TO FIND IT**

Set a specific section as the landing view when 1Password opens, useful for teams that want to land directly on a vault or category.

> 📍 In the app, go to Settings › General › Defaults and select the section to display on launch.

New setting to choose which section of 1Password is shown on first open or after inactivity (Settings > General > Defaults).

Audit event for moving or duplicating items between 1Password accounts now classified as an Export Vault event for clearer audit trails.

`1Password for Mac beta 8.12.26-31 (2026-06-18)` **BLOG ↗** **STABLE**

1Password Mac beta adds configurable default landing section, CTAP2 user-verification support, and improved audit events for item moves.

└──▷ **HOW TO FIND IT**

Set the default landing section so your team's shared vault is the first thing shown each time 1Password opens.

> 📍 In the app, go to Settings › General › Defaults and choose your preferred section from the dropdown.

Adds configurable default landing section so users can choose which part of 1Password is shown on first open or after inactivity (Settings > General > Defaults).

Adds support for CTAP2 authenticators that require user verification during 1Password account sign-in with two-factor authentication enabled.

Changes the audit event logged when an item is moved or duplicated between accounts to an Export Vault event, giving security teams clearer cross-account activity signals.

In Safari, pressing the down arrow in a field now moves focus to the first suggestion in the inline autofill menu.

WAS THIS USEFUL? 👍 👎

Clerk Feedback ⓘ

CLI webhooks, impersonation, and a refreshed experience (2026-07-09)

5 RELEASES · 2026-06-22 → 2026-07-09 BLOG ↗ LATEST STABLE

Clerk is a complete user authentication and management platform that provides embeddable UI components, APIs, and dashboards for adding secure sign-in and identity features to applications.

Clerk CLI 2.0 adds local webhook testing, one-command user impersonation, and a refreshed interactive UX.

New local webhooks toolkit lets you receive and inspect webhook events during development without a public URL.

New one-command user impersonation enables developers to instantly act as any user for debugging and testing.

Refreshed interactive CLI experience redesigned for both human developers and AI agents.

▾ 4 more releases in this issue · 2026-06-22 → 2026-07-09

Disable just-in-time provisioning for SAML connections (2026-07-07) BLOG ↗

STABLE

Clerk SAML connections now support disabling just-in-time provisioning to block unprovisioned users.

└──▷ TRY IT

Disable JIT provisioning on an existing SAML connection so only pre-provisioned users can sign in.

```
$ curl -X PATCH
https://api.clerk.com/v1/saml_connections/<connection_id> \
  -H 'Authorization: Bearer <secret_key>' \
  -H 'Content-Type: application/json' \
  -d '{"disable_jit": true}'
```

Enable or disable JIT provisioning for a SAML connection through the Dashboard without touching the API.

📍 In the Clerk Dashboard, go to Configure › SSO Connections › [your SAML connection] and toggle off Enable just-in-time provisioning.

Adds opt-in control to reject SAML sign-ins from users who lack an existing account, preventing automatic JIT provisioning.

Exposes a `disable_jit` field on the Backend API for programmatic control of JIT provisioning per SAML connection.

Adds a self-serve toggle in the Clerk Dashboard to enable or disable JIT provisioning on SAML connections.

Clerk Billing now supports account credits (2026-06-30) BLOG ↗ STABLE

Clerk Billing adds account credit balance management via Dashboard and API.

Enables adjustment of customer credit balances directly from the Clerk Dashboard.

New API support for programmatically managing customer credit balances in Clerk Billing.

Self-serve SSO (2026-06-26) BLOG ↗ STABLE

Clerk adds self-serve enterprise SSO configuration for IT admins via a new Security tab in OrganizationProfile.

└──▷ **HOW TO FIND IT**

Expose the new Security tab to organization members so their IT admins can self-configure enterprise SSO — no Clerk dashboard access required.

> 📍 In the Clerk dashboard, ensure <OrganizationProfile /> is rendered in your app; IT admins will find the new Security tab there to configure and manage their enterprise SSO connections.

New Security tab in <OrganizationProfile /> lets customers' IT admins configure and manage their own enterprise SSO connections without involving your support team.

Customize your OAuth consent page (2026-06-22) BLOG ↗ STABLE

Clerk lets you host and customize the OAuth consent screen on your own domain via the <OAuthConsent /> component.

└──▷ **USE IT**

Render the OAuth consent screen on your own domain instead of Clerk's hosted page, keeping users in your brand experience.

```jsx
<OAuthConsent />
```

WAS THIS USEFUL?  

Keycloak

Feedback ⓘ Keycloak 26.7.0 released (2026-07-09)

2 RELEASES · 2026-07-03 → 2026-07-09 BLOG ↗ LATEST STABLE

Keycloak is an open-source identity and access management solution that provides single sign-on, authentication, and authorization for applications and services.

Keycloak 26.7.0 adds SCIM API preview, SSF real-time signals, AuthZen PDP, SAML step-up auth, and a new Admin API v2.

└──▷ TRY IT

Enable the SCIM API preview so HR systems or identity governance platforms can provision and deprovision users automatically.

```
$ bin/kc.sh start --features=scim-api
```

Enable real-time SSF security event delivery so downstream apps learn about logouts and credential changes without waiting for token refresh.

```
$ bin/kc.sh start --features=ssf
```

Enable the AuthZen PDP endpoint so PEPs can request standardized permit/deny decisions without coupling to Keycloak's proprietary authorization API.

```
$ bin/kc.sh start --features=authzen
```

Adds SCIM API (preview) for automating user and group provisioning and deprovisioning via standard CRUD, PATCH, filtering, pagination, and schema discovery endpoints.

Adds OpenID Shared Signals Framework (SSF) Transmitter support (experimental), pushing signed Security Event Tokens to relying parties via push (RFC 8935) and poll (RFC 8936) delivery with durable outbox and retry logic.

Adds AuthZen Evaluation API (experimental) implementing the OpenID AuthZEN Authorization API 1.0, exposing single and batch PDP evaluation endpoints for vendor-neutral authorization decisions.

Promotes step-up authentication for SAML clients to supported status, enabling SAML service providers to request specific authentication context classes.

Adds Identity Brokering API v2 with client-level authorization, confidential-clients-only enforcement, OAuth 2.0-compliant POST endpoint, and a new session-based token

storage option.

+6 more

▾ **1 more release in this issue** · 2026-07-03 → 2026-07-09

Experimental Shared Signals Framework support (2026-07-03) BLOG ↗ STABLE

Keycloak gains experimental OpenID Shared Signals Framework (SSF) 1.0 Transmitter support, pushing signed SETs to receivers in seconds.

└──▷ **TRY IT**

Enable the experimental SSF Transmitter feature when starting Keycloak so the per-realm SSF endpoints become available.

```
$ kc.sh start-dev --feature-ssf=enabled
```

Obtain a client-credentials token scoped to SSF management, then use it to create or manage streams against the Keycloak SSF transmitter endpoints.

```
$ curl -s -X POST
"https://my.keycloak.test/realms/myrealm/protocol/openid-
connect/token" \
  -d "grant_type=client_credentials" \
  -d "client_id=caep-receiver" \
  -d "client_secret=…" \
  -d "scope=ssf.read ssf.manage"
```

Enable the SSF Transmitter for a realm and create an SSF-enabled OIDC receiver client, then configure its stream settings via the Admin Console.

📍 In the Admin Console: 1. Open the target realm → Settings → SSF Transmitter toggle → On (sets `ssf.transmitterEnabled`). 2. Go to Clients → Create client, set Client authentication On, Service accounts roles enabled. 3. On the client's SSF tab: enable SSF, set Default Subjects to ALL, add an Audience (e.g. `https://caep.dev`), and tick Push and Poll as supported delivery methods. 4. Under Client scopes, add ssf.read and ssf.manage as Optional scopes.

Adds experimental SSF 1.0 Transmitter role: Keycloak now pushes signed Security Event Tokens (SETs) to subscribed receivers for identity events such as session revocation, credential change, and account disable.

Supports CAEP 1.0 and RISC 1.0 event profiles, automatically mapping native Keycloak events (login, logout, credential change, session revocation) to the correct standardised event types.

> Supports both HTTP push (RFC 8935) and HTTP poll (RFC 8936, return-immediately) delivery channels for SET distribution.
>
> Introduces a transactional outbox: events are durably persisted in the same transaction that produced them, then drained asynchronously with exponential backoff, decoupling generation from delivery.
>
> New synthetic event REST endpoint lets external tools (e.g. SOC/IAM platforms) inject events that did not originate inside Keycloak — useful for bridging external credential-change signals.
>
> +5 more

WAS THIS USEFUL?  

Okta Identity Engine

Feedback ⓘ

Monthly release 2026.07.0 (2026-07-01) | 2 RELEASES · 2026-06-17 → 2026-07-01

BLOG ↗ | LATEST | STABLE

Okta Identity Engine is a flexible authentication platform that enables developers to build customizable, secure identity and access management workflows into their applications.

Okta July 2026 brings bot protection GA, AI agent token improvements, native-to-web SSO, and advanced device posture checks.

└──▷ **TRY IT**

Retrieve only specific fields from the Groups list endpoint to reduce payload size in automated pipelines.

```
$ GET /api/v1/groups?fields=id,profile.name,profile.description
```

Delete a group push mapping stuck in ERROR state to unblock sync without manual console intervention.

```
$ DELETE
/api/v1/apps/{appId}/groups/{groupId}/pushGroups/{pushGroupId}
```

Enable bot protection for a sign-on policy to block automated credential attacks on your org.

> 📍 In the Admin Console, go to Security › Sign-On Policies › select a policy › Edit › enable Bot Protection under the Authentication settings.

Bot protection is now generally available in Production.

New `fields` query parameter for Groups, Realms, and Devices list endpoints is now GA in Production, enabling selective field retrieval.

New IP Service available for enhanced dynamic network zones.

Improved group member search functionality is now GA in Production.

Customizable emails for Passkey (FIDO2 WebAuthn) authenticator is GA in Preview.

+12 more

└──▷ **BREAKING ON UPGRADE**

! The `appId` parameter in the Register an AI agent API is deprecated

! Cross App Access is removed as a self-service feature

▾ **1 more release in this issue** · 2026-06-17 → 2026-07-01

> Weekly release 2026.06.2 (2026-06-17) BLOG ↗ STABLE
>
> Okta Identity Engine adds agent-to-agent connections as Early Access in Preview environments.
>
> Enables agent-to-agent connections (Early Access, Preview) for AI agent API integration scenarios.

WAS THIS USEFUL? 👍 👎

BeyondTrust Feedback ⓘ

Remote Support 25.3.5 release notes (2026-06-30) BLOG ↗

BeyondTrust is a privileged access management platform that secures, controls, and monitors privileged accounts, credentials, and remote access across enterprise environments.

Remote Support 25.3.5 adds LDAP authentication and SCIM integration support through Pathfinder.

Adds LDAP authentication support via Pathfinder, enabling user validation against existing directory services with improved login reliability and error feedback.

Adds SCIM integration for Pathfinder, enabling automated user and access synchronization with identity providers to reduce manual administration.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback

JumpCloud Agent Release Notes for June 16, 2026 (2026-06-16) BLOG ↗

JumpCloud is a cloud-based directory platform that manages user identities, devices, and access to applications and networks across organizations.

JumpCloud Agent 2.138.3 adds JIT URLs for scheduled commands and improves per-user agent reliability on Linux.

Scheduled and repeating commands with remote files now use just-in-time (JIT) URLs, preventing failures when Amazon S3 presigned links expire before execution.

Linux per-user agent more reliably runs during user association and login, enabling JumpCloud Go registration and device trust flows without manual intervention.

WAS THIS USEFUL?

◆ Data Security & DLP

Ubiq Security

Feedback ⓘ snapshot-20260713 NOTES ↗

Ubiq Security is a developer-focused encryption platform that provides APIs and SDKs enabling applications to easily integrate structured data encryption without managing cryptographic complexity.

Ubiq adds Access Groups, dataset tagging, advanced structured encryption config, SIEM streaming, and expanded SDK capabilities.

└──▷ HOW TO FIND IT

Stream Ubiq encryption activity events to Splunk for real-time SIEM monitoring of data access.

> 📍 In the Ubiq Dashboard, go to Administration › Observability Integration Setup and configure your Splunk endpoint to enable Event Streaming.

Enable SSO for the Ubiq Dashboard using your existing enterprise identity provider to centralize access control.

> 📍 In the Ubiq Dashboard, go to Administration › SSO for the Ubiq Dashboard and configure your SAML or OpenID Connect provider (Google Workspace, Microsoft Entra, or Okta).

Monitor Ubiq platform availability within your internal observability stack using the Prometheus healthcheck endpoint.

```
$ curl -H "Authorization: Bearer <auth_token>" https://<ubiq-prometheus-endpoint>/healthcheck
```

Introduces Access Groups to manage dataset access for identities (API Keys and IDP users), replacing the previous Dataset Group → Identity hierarchy.

Adds dataset tagging as a flexible, non-hierarchical replacement for Dataset Groups, allowing grouping by environment, application, or any custom structure.

Enables advanced structured Dataset configuration in the dashboard UI, allowing encryption behavior for dates, numerics, short fields, and large character sets to be defined at policy level without custom code.

Adds long-term event storage with automatic rollups, preserving high-granularity recent data while extending historical visibility to years.

Deploys core system changes with zero downtime to APIs and encryption services.

+18 more

WAS THIS USEFUL? 👍 👎

Immuta

Feedback ⓘ snapshot-20260713 NOTES ↗

Immuta is a data security platform that automates access control, privacy policy enforcement, and data governance across cloud data environments.

Immuta adds automated Redshift discovery, managed policy groups, Alation OAuth 2.0, Delta Sharing masking, and Teams/Slack notifications.

└──▷ **HOW TO FIND IT**

Apply a system-wide subscription policy that restricts AI-restricted data to human users only, using the new built-in All Agents group to exclude agent accounts.

> 📍 In the Immuta console, go to Policies › Subscription Policies › New Policy, set the condition to deny access when the subscriber is a member of All Agents, and scope it to data sources tagged as AI-restricted.

Switch the Alation integration to OAuth 2.0 to eliminate the need for recurring API key rotation every 24–90 days.

> 📍 In the Immuta console, go to App Settings › Integrations › Alation, select OAuth 2.0 as the authentication method, and supply your OAuth client credentials.

Enable automated sensitive data discovery on a Redshift viewless data source to automatically tag columns matching regex, dictionary, or column-name patterns.

> 📍 In the Immuta console, open your Amazon Redshift viewless data source, navigate to the Sensitive Data Discovery tab, and enable automated discovery to begin applying identifier-based tags to matching columns.

Adds automated sensitive data discovery (regex, dictionary, and column name regex identifiers) for the Amazon Redshift viewless integration, auto-tagging matching columns.

Introduces three built-in Immuta-managed groups — All Agents, All Users, and All Identities — usable in subscription and data policies without manual membership management.

Enables case-insensitive Starburst (Trino) principal name matching for more reliable user and service account identification (on by default 7/1, mandatory 8/1).

Revokes data owner grants when a data source is disabled in Immuta (behavior change: opt-in 6/30, on by default 8/1, mandatory 9/1).

Extends Unity Catalog tag ingestion to foreign catalogs (Lakehouse Federation) by leveraging system.information_schema tag views as a fallback.

+7 more

WAS THIS USEFUL?  

DigiCert CertCentral

Feedback ⓘ July 7, 2026

7 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

DigiCert CertCentral is a certificate lifecycle management platform that allows organizations to issue, manage, renew, and automate SSL/TLS and other digital certificates.

CertCentral's Allowed validity periods setting now controls certificate validity (max 199 days), not order validity, with updated TLS request forms.

└──▷ HOW TO FIND IT

Restrict which certificate validity durations your team can select when ordering Basic OV TLS certificates.

📍 In the CertCentral console, go to Settings › Product Settings, select Basic OV in the Product column, configure Allowed validity periods under Product Settings, then click Save Settings.

Allowed validity periods setting for TLS products now controls certificate validity instead of order validity, capped at the 199-day maximum.

Annual plan and Multi-year Plan TLS request forms now show coverage options in years (1-, 2-, 3-year) with custom order validity options removed.

For non-Annual/non-Multi-year Plan TLS requests, certificate validity and order validity are now unified (e.g., a 47-day cert creates a 47-day order).

Breaking change: any previously configured allowed validity periods exceeding 199 days are automatically removed from TLS product settings.

└──▷ BREAKING ON UPGRADE

! The Allowed validity periods setting in TLS certificate product settings no longer controls order validity; it now controls certificate validity. Any previously configured validity periods greater than 199 days will be automatically removed from the setting.

! Customized order validity options have been removed from Annual plan and Multi-year Plan TLS certificate request forms; coverage options now appear only in years (1-, 2-, or 3-year).

! For non-Annual plan and non-Multi-year Plan TLS certificate requests, any validity periods greater than 199 days have been removed, and certificate validity now determines order validity.

▾ 6 more releases in this issue · seen 2026-07-13

June 1, 2026 NOTES ↗ STABLE

CertCentral enforces CT logging for all public TLS certs and upgrades MPIC to four remote locations across two RIRs.

└──▷ **HOW TO FIND IT**

Check whether your existing TLS certificates are logged to CT logs before the enforcement deadline using CertCentral's Report Library.

📍 In CertCentral, go to Reports › Report Library, then generate or run a TLS certificate report and filter by CT log status to identify any certificates currently excluded from CT logs.

Enforces Certificate Transparency (CT) logging for all newly issued public TLS certificates (DV, OV, EV, QWAC, QWAC PSD2) across all DigiCert brands — opt-out is no longer available.

Removes CT logging opt-out options from CertCentral, requiring customers who previously excluded certificates from CT logs to reassess their PKI strategy (private PKI hierarchies remain unaffected).

Upgrades Multi-Perspective Issuance Corroboration (MPIC) to require corroboration from at least four remote network locations spanning at least two different Regional Internet Registry (RIR) regions, aligning with CA/Browser Forum Phase 4 requirements.

└──▷ **BREAKING ON UPGRADE**

! All public TLS certificates (DV, OV, EV, EU QWAC, EU QWAC PSD2) across all DigiCert brands will be logged to CT logs starting June 1, 2026, removing the ability to opt out of CT logging for newly issued, reissued, or duplicate public TLS certificates.

! The CT logging opt-out options are being removed from CertCentral, meaning any workflow or configuration that relied on keeping public TLS certificates out of CT logs will no longer function for certificates issued on or after June 1, 2026.

June 10, 2026 NOTES ↗ STABLE

CertCentral Service Users can now have their first and last names updated via UI or API.

└──▷ **HOW TO FIND IT**

Rename a Service User via the CertCentral console to reflect its operational role or owning team.

📍 In CertCentral, go to Account › Service Users, select a Service User, update the First name and Last name fields under User details, then select Update user.

Enables updating a Service User's first and last name in CertCentral or via the Services API, making it easier to identify API key ownership and assign users to teams or roles.

June 16, 2026 NOTES ↗ STABLE

CertCentral adds domain-scoped API keys, expands Manager org validation permissions, and improves Smart Seal accessibility.

└──▷ **HOW TO FIND IT**

Create a least-privilege API key scoped only to domain management — ideal for automation pipelines that manage domain validation without needing certificate ordering access.

📍 In CertCentral, go to Automation › API Keys › Add API Key. Enter a Description, select the target User, then in the `API key restrictions (optional)` menu select `Domains`, and click Add API key.

New `Domains` API key restriction scope lets CertCentral API keys be limited exclusively to domain management actions (add, submit, validate, revalidate domains) with no access to certificate ordering or account management.

Expands Manager role permissions to include the full organization prevalidation workflow: access the Organizations page, view all organizations, add organizations, and submit them for validation independently of a certificate request.

Improved DigiCert Smart Seal and Site Seal keyboard accessibility: users can now tab to the seal, activate it with Enter or Space, and receive a visible focus ring — giving keyboard-only users the same hover information as mouse users.

Smart Seal and Site Seal now meet WCAG 2.1 criteria (SC 1.1.1, SC 2.1.1, SC 4.1.2): descriptive alt text added to PNG and SVG seal variants, seals identified as interactive buttons to assistive technologies, and screen readers receive the label 'Verify DigiCert TLS certificate for this site'.

June 30, 2026 NOTES ↗ STABLE

CertCentral Service Users now support custom first and last names for clearer API key identification.

└──▷ **HOW TO FIND IT**

When creating a Service User for a specific integration, set meaningful names so admins can instantly identify the key's purpose in the API Keys list.

📍 In CertCentral, go to Automation › API Keys › Add API Key. In the Add API Key window, select `Create new service user` from the User menu, then replace the default `Service` and `User` values in the First Name and Last Name fields (e.g., First Name:

`Integration`, Last Name: `Production`). Complete the remaining fields and select Add API Key.

Enables custom first and last name assignment when creating a Service User, replacing the generic `Service` /'User' defaults to reflect team, integration, or role context.

May 12, 2026 NOTES ↗ STABLE

CertCentral adds editable ACME Directory URL limits and automated authorized representative delegation for PKIo certs in Europe.

:
:

└──▷ **HOW TO FIND IT**

Modify an existing ACME Directory URL's allowed domains and order cap instead of creating a new one — useful when onboarding new subdomains or scaling automation quotas.

📍 In CertCentral, go to Automation › ACME Directory URLs. In the Actions menu for the target URL, select Update ACME Directory URL. Under Allowed SANs (optional), add or remove domains. Under Number of allowed orders (optional), increase the limit or remove it to restore unlimited status. Select Update ACME Directory URL.

Set up authorized representative delegation for a service user so PKIo Private Services Server certificate orders are auto-approved without repeated manual sign-offs.

📍 In CertCentral Europe, go to Account › Users. Select the target user. Under Authorized representative delegation, select Request authority delegation. In the Submit a delegation request window, choose the organization and authorized representative, then select Submit.

Enables editing of domain (SAN) and order limit restrictions on existing ACME Directory URLs — add/remove allowed SANs, adjust order caps, or restore unlimited status without creating a new URL.

Introduces automated Authorized Representative Delegation in CertCentral Europe for PKIo Private Services Server certificates, allowing PKIo orders to be auto-approved after a one-time delegate setup — eliminating per-order manual approvals.

Supports delegation of authorized representative authority to administrators, managers, and service users, enabling bulk certificate issuance via CertCentral Europe or the Services API.

September 15, 2026 NOTES ↗ STABLE

DigiCert CertCentral now issues public TLS certificates with partitioned CRL Distribution Point URLs, capping each CRL at ~9.5 MB.

> All newly issued public DigiCert TLS certificates include CRLDP URLs pointing to partitioned CRLs (~9.5 MB each) instead of a single monolithic revocation list.

WAS THIS USEFUL? 👍 👎

Cyera

Feedback ⓘ AI-powered Issue Severity

15 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Cyera is a cloud data security platform that discovers, classifies, and protects sensitive data across cloud environments to reduce risk and ensure compliance.

Cyera adds AI-Powered Issue Severity to evaluate alerts in full business context — data sensitivity, identity access, and exposure conditions.

New AI-Powered Severity engine scores issues using business context including data sensitivity, identity access, and exposure conditions to prioritize what truly matters.

Provides AI-generated explanations for each severity rating to accelerate confident remediation decisions.

▾ **14 more releases in this issue** · seen 2026-07-13

> Access Trail NOTES ↗ STABLE
>
> Cyera launches Access Trail: near real-time sensitive data activity monitoring with one year of retention.
>
> New Access Trail capability captures near real-time activity tied to sensitive data, recording access, movement, and exposure over time.
>
> Supports one year of activity history retention for incident investigations, access pattern validation, and audit support.

> Automation Center NOTES ↗ STABLE
>
> Cyera launches Automation Center: a centralized hub to turn data security insights into consistent, repeatable remediation workflows.
>
> New Automation Center module enables teams to define, scope, and execute repeatable data security remediation workflows across the Cyera platform.
>
> Supports automated delegation of data ownership directly from workflow definitions.
>
> Enables condition-based enforcement of security controls without manual coordination, scaling consistently as environments grow.

> Browser Shield NOTES ↗ STABLE

Cyera Browser Shield adds real-time AI prompt monitoring and inline blocking to control Shadow AI across the enterprise.

.

└──▷ **HOW TO FIND IT**

Enable Browser Shield visibility and inline blocking through the Cyera console to surface and control Shadow AI usage across your organization.

> 📍 In the Cyera console, navigate to Browser Shield, enable monitoring for browser-based AI tools, configure policies to distinguish corporate vs. personal accounts, and connect alert output to Omni DLP for unified investigation.

New Browser Shield capability provides enterprise-wide visibility into approved and unsanctioned browser-based AI tool usage.

Monitors AI prompts in real time before submission, enabling interception of risky inputs before data leaves the browser.

Distinguishes between corporate and personal accounts on AI tools, enabling policy enforcement based on account context.

Blocks risky prompt submissions inline, preventing sensitive data from reaching unsanctioned AI services.

Integrates Browser Shield alerts with Omni DLP for unified investigation of AI-related data exposure events.

Cyera MCP NOTES ↗ STABLE

Cyera MCP connects any MCP-compatible AI tool to DataPort for natural-language data security queries and agent workflows.

.

New Cyera MCP integration enables building custom data security agents by connecting MCP-compatible AI tools to Cyera's DataPort.

Supports natural-language querying of sensitive data locations, access paths, exposure risks, and classification insights via DataPort.

Enables AI-driven threat hunting, investigation acceleration, and report generation grounded in live Cyera security intelligence.

Cyera Privacy NOTES ↗ STABLE

Cyera Privacy launches with continuous data discovery, RoPA, DSR automation, risk management, and consent management.

.

New continuous data discovery and enriched classification engine grounds privacy operations in real-time inventory of personal data locations and usage.

Adds automated data mapping and Records of Processing Activities (RoPA) built from live data context rather than manual surveys.

Introduces privacy assessments pre-filled with actual data discovery findings, reducing manual effort in DPIA and similar workflows.

Adds Data Subject Request (DSR) automation to fulfill access, deletion, and portability requests against a continuously updated data inventory.

New centralized privacy risk management surface consolidates privacy risks identified across the data estate.

+1 more

DLP Trends NOTES ↗ STABLE

Cyera DLP Trends surfaces recurring exposure patterns across users, destinations, and roles instead of single-event alerts.

.

New DLP Trends view aggregates repeated exposure events into patterns, showing what keeps happening, where, and who is involved.

Enables prioritization of recurring risks — such as data repeatedly sent to personal destinations or sharing tied to specific roles — over one-off alert triage.

Incorporates Cyera's data context to elevate provider severity scores within trend analysis.

Data Lineage for Files NOTES ↗ STABLE

Cyera adds Data Lineage for Files, tracing how humans and AI move sensitive files across system boundaries.

.

New Data Lineage for Files capability traces the full lifecycle of sensitive files as they are moved, transformed, and propagated by humans and AI agents across systems.

Enables reconstruction of file propagation paths even when files cross system boundaries.

Turns isolated security events into connected investigation context within a single lineage view.

Data Retention Policies NOTES ↗ STABLE

Cyera adds Data Retention Policies to surface overdue and stale sensitive files at scale across cloud environments.

.

New Data Retention Policies capability detects files that have exceeded a defined

retention period using file age signals combined with Cyera's classification context.

Identifies sensitive data no longer in active use, enabling teams to reduce audit scope and breach impact without manual review at scale.

Adds visibility into existing Microsoft Purview retention labels applied to M365 files, supporting audit readiness and compliance assessment for Microsoft customers.

Exchange Online Coverage NOTES ↗ STABLE

Cyera adds Exchange Online coverage, scanning mailboxes, email content, and attachments for sensitive data.

.

Supports Exchange Online as a new data source, enabling discovery and classification of mailbox content, email bodies, and attachments.

Reveals regulated and sensitive data spread across historical email communications to support cleanup and retention enforcement.

Provides visibility into how sensitive attachments move through email and whether email is used beyond its intended governance scope.

Identity Visibility for On-prem NOTES ↗ STABLE

Cyera extends identity visibility to on-prem file shares via Active Directory mapping across SMB, NetApp, and Dell PowerScale.

.

Adds identity visibility for on-prem file shares by mapping Active Directory access across SMB File Shares, NetApp, and Dell PowerScale.

Enables teams to see who has access, how permissions were granted, and where sensitive data is overly exposed in on-prem environments.

Surfaces access paths (rather than individual files) to help teams identify and reduce overly permissive or inherited access at the source.

Recent Integrations NOTES ↗ STABLE

Cyera adds nine new integrations spanning cloud security, NHI, data catalogs, lineage, and vuln prioritization.

.

New integration with AWS Security Hub to surface Cyera data risk context alongside cloud security posture findings.

New integration with Upwind to enrich cloud security workflows with data exposure context.

New integration with Clutch to bring data risk context into non-human identity remediation workflows.

New integration with Atlan to extend data risk and sensitivity context into data catalog workflows.

New integration with Euno to incorporate data lineage context into risk triage.

+4 more

ServiceNow ITSM Coverage NOTES ↗ STABLE

Cyera adds ServiceNow ITSM as a discoverable data source, scanning tickets, incidents, attachments, and free-text fields.

.

Adds ServiceNow ITSM as a supported data source, enabling discovery and classification across tickets, incident records, attachments, and free-text fields.

Detects exposed PII, plain-text credentials, and regulated data within day-to-day ITSM operational workflows that traditional tools overlook.

Brings ServiceNow under Cyera's DSPM governance, treating it as a managed data store rather than an unmonitored exception.

Slack Integration NOTES ↗ STABLE

Cyera adds Slack integration to deliver actionable data security alerts directly to data owners.

└──▷ HOW TO FIND IT

Route data security alerts to the relevant Slack channel or owner so issues are surfaced and acted on where teams already work.

📍 In the Cyera console, go to Integrations › Slack and connect your Slack workspace to enable alert delivery to data owners.

New Slack integration sends actionable data security alerts with context and next steps directly to data owners in Slack.

Enables parallel remediation workflows so data owners can act without routing through security teams.

Topic Exploration NOTES ↗ STABLE

Cyera Topic Exploration groups technical data classifications into business-aligned concepts for clearer policy and faster remediation.

.

Introduces Topic Exploration, which aggregates multiple technical data classifications into business-aligned topics that map to how organizations think about their data.

> Enables policy creation, automation, and remediation workflows to be driven by business context rather than raw data labels.
>
> Allows security and data teams to prioritize and communicate data risk in terms that business leaders understand.

WAS THIS USEFUL? 👍 👎

Netskope

Feedback ⓘ

Netskope Cloud 131.0.0 Documentation Available Now!

8 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Netskope is a cloud security platform that provides real-time visibility, threat protection, and data security across cloud apps, websites, and private applications.

Netskope R131 adds Next-Gen SSPM onboarding for M365/Salesforce, on-demand deep scanning, on-prem detection, and macOS/iOS 26 client support.

└──▷ HOW TO FIND IT

Onboard a Microsoft 365 tenant into Next-Gen SSPM to continuously assess your M365 security posture against benchmarks.

> 📍 In the Netskope console, go to SSPM › Next-Gen SSPM › Onboarding and select Microsoft 365, then follow the guided connector setup.

Adds pre-built Ubuntu 22.04 VM image for Hyper-V to simplify Netskope Private Access publisher deployments.

Supports onboarding Microsoft 365 in Next-Gen SSPM for SaaS security posture visibility.

Supports onboarding Salesforce in Next-Gen SSPM for SaaS security posture visibility.

New Deepscan On Demand capability for Threat Protection API enables targeted deep file inspection outside of inline policy.

New On-Prem Detection for Traffic Steering identifies and steers on-premises traffic.

+4 more

▾ 7 more releases in this issue · seen 2026-07-13

> Netskope Cloud 132.0.0 Documentation Available Now! NOTES ↗ STABLE
>
> Netskope R132 adds Universal Reverse Proxy for CASB, Transaction Event visibility in Skope IT, and enhanced SSPM Findings UI.
>
> Adds Universal Reverse Proxy support under CASB Real-time Protection.
>
> New Transaction Event support in Skope IT for SaaS visibility.
>
> Enhanced Findings UI in SaaS Security Posture Management (SSPM).
>
> New visibility for Modified and Deleted Rules in the SSPM Rules UI.

> Netskope Cloud 133.0.0 Documentation Available Now! NOTES ↗ STABLE

Netskope R133 adds AI Red Teaming, AI Security Guardrails, an AI Gateway, expanded ChatGPT Enterprise protection, and NPA source-IP egress controls.

New AI Red Teaming capability to test and evaluate AI model resilience against adversarial prompts and attacks.

New AI Security Guardrails to enforce policy-based controls on AI application usage and outputs.

New Netskope AI Gateway to provide a centralized security enforcement point for AI service traffic.

Expanded Next Generation API Data Protection for ChatGPT Enterprise, covering fine-tuning data and memory data.

Controlled GA of Windows 64-bit Client Support Mandate for the Netskope Client.

+4 more

Netskope Cloud 134.0.0 Documentation Available Now! NOTES ↗ STABLE

Netskope R134 adds Enterprise Browser–NPA integration, GitHub 3rd-party app risk scoring, and a new Threat Hunting Module.

New Next Generation API Data Protection deferred policy action capability.

New Enterprise Browser and NPA (Network Private Access) integration.

New SSPM capability to discover and risk-score third-party apps in GitHub.

New deep visibility into SaaS users within SaaS Security Posture Management.

New Threat Hunting Module added to Threat Protection.

Netskope Cloud 135.0.0 Documentation Available Now! NOTES ↗ STABLE

Netskope R135 adds AI Guardrails, AI Red Teaming, Agentic Broker, and new CASB/SSPM integrations.

New AI Prompt and Response Guardrails for content moderation and protection from AI-based threats.

New AI Red Team product for adversarial testing of AI systems.

New Agentic Broker capability in CASB Real-time Protection.

General Availability of Smartsheet support in Next Generation API Data Protection.

New GitHub Activity Alerts and Dashboard support in Next Generation API Data Protection.

+1 more

Netskope Cloud 138.0.0 Documentation Available Now! NOTES ↗ STABLE

Netskope R138.0.0 adds enhanced selection controls for Golden Monthly Releases in Client Steering.

Netskope Cloud 139.0.0 Documentation Available Now! NOTES ↗ STABLE

Netskope R139 adds FIPS-140-3 client steering, device deletion, and GA of AI Command Center.

Adds FIPS-140-3 compliance support for Netskope Client Steering traffic.

New ability to delete devices from the Netskope management console.

AI Command Center (AICC) reaches General Availability (GA).

The Netskope Cloud Platform R137.0.0 documentation is here! NOTES ↗ STABLE

Netskope R137 adds SSPM Privileged User Visibility enhancements and a self-service UI for NPA AI Agent App Audit.

Adds Privileged User Visibility enhancements to SaaS Security Posture Management (SSPM), surfacing elevated-access account risk in SaaS environments.

Introduces a self-service UI for the NPA AI Agent App Audit, enabling administrators to review and manage AI agent application access without requiring support intervention.

WAS THIS USEFUL? 👍 👎

Varonis

Feedback ⓘ Access request shortcut in Windows Explorer

10 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Varonis is a data security platform that monitors, protects, and manages sensitive data across cloud and on-premises environments to prevent insider threats and cyberattacks.

Varonis adds a Windows Explorer right-click shortcut so users can request file access without IT intervention.

└──▷ HOW TO FIND IT

Grant end-users a self-service access request flow from within Windows Explorer to reduce IT helpdesk tickets.

📍 In Windows Explorer, right-click a file or folder you lack access to, then select the Varonis access request option from the context menu to submit a request.

New right-click context menu in Windows Explorer lets users request access to files and folders directly, without copying paths or contacting IT.

▾ 9 more releases in this issue · seen 2026-07-13

Centralized MDDR escalation management NOTES ↗ STABLE

Varonis MDDR dashboard gains a direct link to Varonis Community for centralized escalation contact management.

└──▷ HOW TO FIND IT

Reach escalation contact management without leaving the MDDR dashboard — navigate directly from the new link instead of hunting through Community separately.

📍 In the Varonis Web Interface, go to the MDDR dashboard and click the Varonis Community referral link to manage escalation contacts.

Configurable rescan behavior for classification policies NOTES ↗ STABLE

Varonis now lets admins choose between incremental and full rescans when editing classification policies.

└──▷ HOW TO FIND IT

When updating a sensitive-data classification policy in a large environment, choose incremental mode to avoid a costly full rescan and limit impact to only new and modified files.

📍 In the Varonis console, navigate to your classification policy, make the desired edits, and when prompted select 'Apply to new and modified files only' or Full rescan of

all files.

Adds configurable rescan behavior for classification policies: apply edits to new/modified files only, or trigger a full rescan of all files.

Exclude groups from request-based access workflows NOTES ↗ STABLE

Varonis lets admins exclude groups from request-based access workflows via the DAG API.

Group permission updates in near real-time NOTES ↗ STABLE

Varonis now propagates group-based permission changes in near real-time, eliminating access management delays.

Delivers group-based permission updates in near real-time, reducing lag when modifying access group memberships or entitlements.

Localized timestamps for Data Access Governance requests NOTES ↗ STABLE

Varonis Data Access Governance request timestamps now display in the customer's local time zone.

NFS support for Hitachi NAS NOTES ↗ STABLE

Varonis adds NFS monitoring for Hitachi NAS, covering activity auditing, permissions tracking, and data classification.

Supports NFS monitoring for Hitachi NAS data sources, enabling file and folder activity auditing.

Tracks permissions and ownership changes on Hitachi NAS over NFS.

Enables data classification and labeling for Hitachi NAS NFS content.

Prioritized folder rescanning NOTES ↗ STABLE

Varonis adds on-demand prioritized folder rescanning to accelerate classification validation in critical areas.

└──▷ HOW TO FIND IT

After updating a sensitive-data policy, trigger an immediate rescan of the folders most likely to be affected rather than waiting for the next scheduled full scan.

📍 In the Varonis console, navigate to the folder(s) you want to prioritize, trigger an on-demand scan, and set the scan priority order to place those folders first.

Enables on-demand folder rescans with user-defined priority order, so high-value or policy-sensitive folders return updated classification results first.

Review and enforce automatic governance rules NOTES ↗ STABLE

Varonis DAG API now exposes automatic governance rules for governed groups, enabling programmatic visibility and on/off control.

└──▷ TRY IT

Query the DAG API to audit which automatic governance rules are active on a governed group before enforcing a policy change.

```
$ curl -X GET https://<varonis-host>/api/dag/v1/governed-
groups/<group-id>/automatic-rules \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json'
```

Disable an automatic governance rule on a governed group during a planned access review window.

```
$ curl -X PATCH https://<varonis-host>/api/dag/v1/governed-
groups/<group-id>/automatic-rules/<rule-id> \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"enabled": false}'
```

New DAG API endpoints provide full visibility into automatic governance rules configured on governed groups.

Enables programmatic enforcement: turn automatic governance rules on or off on demand via the DAG API.

Secure Cursor with Varonis Atlas NOTES ↗ STABLE

Varonis Atlas now secures Cursor AI coding agents with real-time policy enforcement, session logging, and MCP skill inspection.

Adds direct integration with Cursor, the AI-native coding tool, providing real-time visibility into agent intent, data sensitivity, access patterns, and token usage.

Enables runtime policies to block, modify, or alert on risky agent actions before they impact sensitive data.

Introduces detailed session logging capturing prompts, tool usage, and executed actions for incident investigation and response.

> Extends visibility to Cursor configurations and artifacts across runtime environments and source repositories including GitHub and Bitbucket.
>
> Detects malicious MCP skills that attempt to manipulate agents, exfiltrate data, or introduce unsafe actions.

WAS THIS USEFUL?  

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Secureframe

Feedback ⓘ June 2026

13 RELEASES · 2025-03-01 → 2026-06-01 NOTES ↗ STABLE

Secureframe Product Updates is a changelog platform that communicates new features, improvements, and updates to Secureframe's compliance and security automation software.

Secureframe June 2026 adds multi-policy test linking, reopenable reviews, and daily 81–98% faster personnel syncs.

Supports linking multiple policies to a single test, with draft-policy status surfaced as the failure reason.

Enables automatic re-running of paired acknowledgement tests when a policy is published, eliminating stale results.

Adds ability to reopen completed reviews or individual applications to correct mistakes without losing the audit trail.

Enables removal of specific test runs or errors from the Complete tab while preserving audit history.

Personnel and access data now syncs daily using incremental pulls, reducing sync times by 81–98% and lowering Microsoft Graph API quota pressure — no configuration required.

▾ 12 more releases in this issue · 2025-03-01 → 2026-06-01

May 2026 NOTES ↗ STABLE

Secureframe May 2026 adds CMMC Scoping Module, Defense Navigator, GCC High automation, AI-populated SSPs, and overhauled evidence management.

└──▷ HOW TO FIND IT

Keep employee acknowledgment records current immediately after saving a revised policy, without extra manual steps.

📍 In the console, go to Policies › [select policy] › Edit, make your changes, and enable Reset existing acceptances before clicking Save. Employees are automatically notified to re-acknowledge.

Reduce policy maintenance overhead across multiple templates by conditionally showing or hiding sections based on which frameworks or integrations apply.

📍 In the console, go to Policies › [select policy] › Edit, insert a Conditional Token around the relevant section, and define the condition (e.g., framework = CMMC or integration = GCC High) so the section renders only when that condition is met.

New Scoping Module guides organizations through CMMC scoping questions, auto-marks out-of-scope requirements with audit-ready justifications, disables irrelevant tests, and tailors SSP and policy language accordingly.

New Microsoft GCC High integration automates CMMC-required configurations and provides step-by-step manual guidance where Microsoft APIs don't permit full automation.

SSP creation now auto-populates implementation statements from scoping decisions and connected integrations, with bulk or individual acceptance of suggested updates.

CMMC Level 2 framework updated with revised control mappings, refreshed policies, and new automated tests for GCC High, Google Workspace, Federal MDM, and Virtual Desktop environments.

Windows 11 virtual desktops now ship with Microsoft 365 pre-installed and support Common Access Card (CAC) authentication, with automatic smart card redirection and DoD certificate chain configuration.

+4 more

April 2026 NOTES ↗ STABLE

Secureframe April 2026 adds a faster Org Picker with role/team visibility and improved Personnel access tables.

└──▷ **HOW TO FIND IT**

Quickly validate which systems and devices are tied to a specific employee during an audit preparation.

📍 In the console, go to Personnel › Accounts or Personnel › Devices, then use the updated sort and filter controls to narrow results by user.

Switch between client or team organizations without losing context on your role or needing to re-authenticate.

📍 In the console, open the Org Picker, review the role and team size displayed per org, pin frequently used orgs to Favorites, then select the target org or use the built-in sign-out option.

New Org Picker shows role and team size per organization at a glance, with an improved favorites section and built-in sign-out access for faster multi-org

navigation.

Updated Accounts and Devices tables in Personnel add improved sorting, filtering, and navigation to streamline user access and device assignment reviews.

March 2026 NOTES ↗ STABLE

Secureframe adds GA User Access Reviews, guided ComplyAI risk workflows, a dedicated Auditors view, and flexible user onboarding controls.

└──▷ HOW TO FIND IT

Launch a recurring User Access Review for an application to generate audit-ready evidence for SOC 2 or ISO 27001 — use CSV upload to cover apps without native integrations.

📍 In Secureframe, go to Comply › User Access Reviews, click New Review, select Recurring, choose the target application, upload a CSV for non-integrated apps, assign reviewers, and click Launch Review.

Adds User Access Reviews (UAR) — generally available for Comply Complete customers — enabling scheduled or one-time access reviews with approve/revoke/follow-up decisions, bulk actions, CSV uploads for non-integrated apps, and exportable audit-ready evidence across SOC 2, ISO 27001, and CMMC.

Improves ComplyAI risk assessment workflow with guided field completion (risk description, owner) before AI recommendations are generated, surfacing the next required step dynamically.

Introduces a dedicated Auditors tab within Personnel, separating auditor accounts into a tailored view aligned to their role.

Enables admins to control whether users without a start date are automatically invited, allowing onboarding workflows to match organizational timelines.

February 2026 NOTES ↗ STABLE

Secureframe adds Smart User Mentions, Audit Completion Workflow, multilingual onboarding, and expanded Azure/Sentinel integrations.

└──▷ HOW TO FIND IT

Keep audit feedback inside Secureframe instead of Slack — mention your auditor directly on a failing test so they receive an email with the comment and a deep link to that exact test.

📍 In the Audit Module, open the relevant test, add a comment, and type '@' followed by the auditor's name to trigger an email notification with a direct link back to the test

and comment.

Formally close out an audit lifecycle by uploading all report versions and confirming receipt of the final report without leaving Secureframe.

> 📍 In the Audit Module, open the audit and navigate to the Audit Completion modal. Upload draft, revised, and final report files, use the built-in comments with @mentions for feedback, then enable the `final report received` toggle and mark the audit as complete.

Adds Smart User Mentions: tag customers or auditors with '@' in comments across the platform, triggering email notifications with a direct link back to the relevant comment and test location in the Audit Module.

Introduces a structured Audit Completion Workflow with a dedicated modal for uploading multiple report versions (drafts, revisions, finals), a built-in comments section with @mentions, a toggle to confirm final report receipt, and the ability for either party to mark the audit complete.

Supports employee onboarding workflows — including policy acknowledgments and training steps — in French, German, and Spanish for global teams.

Expands Azure integration to ingest user data, increasing identity-based compliance testing coverage in the Access tab.

Adds support for the government version of Microsoft Sentinel, extending SIEM coverage for regulated and public sector environments.

+2 more

January 2026 NOTES ↗ STABLE

Secureframe adds new audit test statuses, 100% Met progress tracking, log deep-links, and progressive auditor-firm search.

.
.
.

└──▷ **HOW TO FIND IT**

Link your audit firm without manual lookup by typing its name in the progressive search field during auditor setup.

> 📍 In the platform, go to the Audit Module › Auditor Setup, type your firm's name in the search field, and select the matching firm from the Auditor Partner Console (APC) results.

Quickly navigate to a specific change during an active audit without scrolling through the full audit trail.

📍 In the Audit Module, open the Audit Log table, locate the relevant entry, and click the direct link in the row to jump to the changed item.

New test statuses in the Audit Module improve transparency between auditors and admins throughout the audit lifecycle.

Audit Module progress bar now tracks completion toward 100% Met, giving an at-a-glance readiness indicator.

Audit Log gains direct links to changed items, enabling users to jump instantly from the log table to any updated record.

Progressive search for audit firms lets users type a firm name to search across all partners with an Auditor Partner Console (APC), enabling self-service auditor linking.

December 2025 NOTES ↗ STABLE

Secureframe December 2025 adds a full Audits Module, FedRAMP 20x JSON export, Trust Center versioning, and automated test-status email alerts.

.
.
.
.

└──▷ **HOW TO FIND IT**

Keep external stakeholders current when a security report or certification is renewed — no manual outreach needed.

📍 In the Trust Center, open an existing resource, select `Replace Resource`, upload the new version, and enable `Notify previously approved users` before saving.

Give non-admin users visibility into only the tests they own so they can act on emerging failures between audits.

📍 In Secureframe, go to Settings › Notifications, set `Test Status Change` frequency to `Weekly`, and confirm the recipient scope — Admins/Super Admins see all tests; other users see only tests they own.

Produce a structured FedRAMP 20x KSI submission artifact from current control and evidence data for assessor review.

📍 In the Secureframe dashboard, navigate to Reports › FedRAMP 20x and select `Export KSI JSON` to download the machine-readable file containing system details, control health, test results, and evidence references.

New Audits Module centralizes the entire audit lifecycle — admins create audits, invite auditors, assign internal stakeholders, and a complete audit history is stored for future reference.

Auditors can review requests, accept or reject evidence, request additional information, and add new requests directly within the platform.

Evidence uploads now capture an Activity Completion Date, distinguishing when work occurred from when evidence was submitted to improve audit scoping accuracy.

New FedRAMP 20x machine-readable JSON export aligns to Key Security Indicators (KSIs), including system details, control health, test results, and evidence references for assessor review.

Trust Center gains a resource-versioning workflow: teams can replace an existing document and automatically notify previously approved users when a new version is available.

+1 more

October 2025 NOTES ↗ STABLE

Secureframe adds Gap Assessments, Snowflake & Freshservice integrations, Custom Integrations, GovRAMP Core, and bulk employee scoping.

.

└──▷ **HOW TO FIND IT**

Evaluate readiness against a specific framework before an audit to surface and prioritize control gaps in one view.

> 📍 In the Secureframe console, navigate to Frameworks › select your target framework (e.g., ISO 27001) › click Run Gap Assessment to generate a prioritized list of missing controls.

Route compliance tasks to your internal IT service desk and automatically track SLA adherence for vulnerability and access review tickets.

> 📍 In the Secureframe console, go to Integrations › ITSM › Freshservice › Connect, then set Freshservice as the task destination for the desired compliance workflow (e.g., Vulnerability Management) to enable bi-directional ticket sync.

Automate evidence collection from a homegrown or unsupported system by mapping its resources into Secureframe's compliance engine.

> 📍 In the Secureframe console, go to Integrations › Custom Integrations › New Integration, then define resource types (e.g., user accounts, devices) and upload or connect your data source to begin automated compliance testing.

Adds Gap Assessments for any supported framework (SOC 2, ISO 27001, NIST 800-171, and more), automatically comparing current posture against requirements and

surfacing prioritized remediation actions.

New Snowflake integration automates continuous compliance testing across authentication, high availability, firewall rules, and alerting for the cloud data platform.

New Freshservice integration enables bi-directional ticket syncing as a task destination to track compliance for incidents, vulnerabilities, nonconformities, system changes, and user access against SLAs.

Introduces Custom Integrations, allowing teams to ingest, map, and test compliance resources (user accounts, devices, training records, cloud infrastructure) from any system not covered by native integrations.

Adds GovRAMP Core framework support, a streamlined baseline derived from GovRAMP Low and Moderate to simplify the authorization process.

+1 more

September 2025 NOTES ↗ STABLE

Secureframe adds EU AI Act compliance framework, MongoDB Atlas integration, and Freshdesk bi-directional sync

Adds full EU Artificial Intelligence Act (EU AI Act) compliance framework with a risk-based approach to categorizing and safeguarding AI systems.

New MongoDB Atlas integration enables automated compliance testing for authentication, backup and availability, encryption, firewall rules, and logging and alerting.

New Freshdesk integration with bi-directional syncing automates evidence collection for customer support and ticketing workflows, eliminating manual uploads for audit readiness.

August 2025 NOTES ↗ STABLE

Secureframe adds 7 new vuln-scanner integrations, Addigy MDM support, C5 compliance, and overhauled vendor review scheduling.

.
.
.

└──▷ HOW TO FIND IT

Clean up legacy policies quickly rather than archiving them one by one before an audit.

> 📍 In Secureframe, go to Policies, select the checkboxes for the outdated policies you want to remove, then choose Archive from the bulk action bar.

Centralizes vendor review schedules so frequency, question sets, requested documents, and vendor lists are configured at the review level rather than per

vendor.

Adds `requested documents` to review schedules, automatically prompting vendors to upload items like SOC 2 reports and spinning up a vendor portal on the review start date.

New integrations with Rapid7 InsightVM, Tenable VMS, Qualys, Wiz, Microsoft Defender, Aikido, and Snyk automate evidence collection for vulnerability scanning, dependency testing, and SAST.

New integration with Addigy enables automated compliance testing for hard drive encryption, firewall configuration, screen lock enforcement, and anti-malware on macOS devices.

Supports the C5 Controls Catalog (BSI) for organizations operating in the EU or targeting the German public sector.

+5 more

June 2025 NOTES ↗ STABLE

Secureframe launches MCP Server, HubSpot integration, self-service SSO, and a refreshed sidebar in June 2025.

.

└──▷ **HOW TO FIND IT**

Ask an AI assistant to surface failing SOC 2 controls directly from Secureframe without logging into the platform.

> 📍 In Claude or Cursor (with the Secureframe MCP Server connected), prompt: "Show me all failing controls for our SOC 2 program in Secureframe."

Enable self-service SSO and domain verification to enforce login restrictions without a support ticket.

> 📍 In the Secureframe console, go to Settings › Domain Management, claim and verify your domain, then enable SSO and configure login restrictions from the same self-service flow.

New Secureframe MCP Server (public beta): a read-only Model Context Protocol server lets AI assistants (ChatGPT, Claude, Cursor) query compliance data — failing controls, high-risk vendors, framework status across SOC 2, ISO 27001, CMMC, and FedRAMP — using natural language in real time.

New HubSpot integration: scans HubSpot users and access data to support compliance-related requirements, centralizing CRM data into the compliance program.

New self-service domain and SSO management: customers can claim and verify domain ownership, enable single sign-on, and configure login restrictions entirely

within the platform without requiring vendor assistance.

Refreshed sidebar navigation: users can now switch between modules even when the sidebar is collapsed, with a cleaner and more consistent layout.

May 2025 NOTES ↗ STABLE

Secureframe May 2025 adds AI Evidence Validation, FedRAMP 20x support, Email Log, and new Vercel/Supabase integrations.

Adds AI Evidence Validation that analyzes uploaded evidence files and metadata in real time, checking document relevance to the control under test and whether timestamps fall within the required testing window.

Supports the FedRAMP 20x KSI framework out of the box, with Phase One pilot key security indicators pre-mapped to controls and tests for FedRAMP Low KSI requirements — part of the FedRAMP 20x pilot alongside Coalfire.

New OneSchema-powered device data upload experience provides instant template validation feedback and highlights file issues before submission, with backend performance improvements for faster, more reliable uploads.

New centralized Email Log gives admins visibility into all outbound Secureframe communications (onboarding invites, task notifications, role changes) with searchable, filterable delivery status and open history for messages sent on or after March 17, 2025.

Revamped User Access Review Access Table built on Castle components with improved navigation, filtering, and search capabilities.

+2 more

March 2025 NOTES ↗ STABLE

Secureframe March 2025 adds an Audit Module, Zendesk integration, CMMC personnel fields, and new auth policies.

└──▷ HOW TO FIND IT

Use Zendesk as a compliance task destination so ticket updates flow into Secureframe and map to your compliance tests — ideal for teams already triaging incidents or vulnerabilities in Zendesk.

📍 In the Secureframe console, go to Settings › Integrations, select Zendesk, and enable bi-directional sync. Choose the record types to import (incidents,

nonconformities, vulnerabilities, system changes, user access) and map them to your compliance tests.

Streamline vendor security reviews by uploading existing questionnaire documents rather than re-entering data manually.

📍 In the Secureframe console, navigate to Security Questionnaires, click Import, and upload a .doc or PDF file to populate the questionnaire module.

Introduces a new Audit Module designed to streamline the end-to-end audit process for customers and auditors.

Refreshes the Dashboard with an updated design and a new Action Items widget surfacing expiring evidence, upcoming reviews, and task due dates.

Adds bi-directional Zendesk integration for task syncing and ticket import, covering incidents, nonconformities, vulnerabilities, system changes, and user access mapped to compliance tests.

Launches a new Help Flyout replacing the old help menu, with direct access to Secureframe Academy, the Get Started page, and Customer Success.

Adds two new CMMC-focused columns to the Personnel Table to support Access Control and Account Management requirements.

+4 more

WAS THIS USEFUL? 👍 👎

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.32.1.0 (2026-07-08)

7 RELEASES · 2026-06-16 → 2026-07-08 BLOG ↗ LATEST STABLE

RegScale Changelog is a feed tracking updates, new features, bug fixes, and improvements to the RegScale compliance and security management platform.

RegScale 6.32.1.0 adds PIV card authentication architecture for high-assurance identity environments.

Supports PIV card authentication architecture, enabling organizations with advanced identity requirements (e.g., federal agencies) to integrate Personal Identity Verification-based access into the platform.

▾ 6 more releases in this issue · 2026-06-16 → 2026-07-08

RegScale 6.32.0.0 (2026-07-03) BLOG ↗ STABLE

RegScale 6.32.0.0 introduces a redesigned UI navigation experience with new features available exclusively in it.

New UI navigation experience replaces the legacy layout; most new features and improvements are available only in the new UI going forward.

RegScale 6.31.1.4 (2026-06-29) BLOG ↗ STABLE

RegScale 6.31.1.4 adds Wiz STIG sync via XCCDF Benchmark and pre-flight validation for eMASS exports.

Supports synchronizing Wiz STIG checks through the XCCDF Benchmark workflow for streamlined Wiz integration.

Introduces pre-flight validation for eMASS exports to surface potential issues before submission.

RegScale 6.31.1.3 (2026-06-25) BLOG ↗ STABLE

RegScale 6.31.1.3 adds XCCDF TestResults import and enhanced CKL/CKLB ingestion for broader assessment data coverage.

Supports importing XCCDF TestResults into the assessment pipeline, enabling additional security assessment data sources within RegScale.

Enhanced CKL/CKLB import capabilities for expanded vulnerability management and assessment data ingestion.

CLI 6.35.0 (2026-06-23) BLOG ↗ STABLE

RegScale CLI 6.35.0 streams Nessus imports and routes STIG/CIS compliance findings to checklists and POA&Ms.

Nessus scan imports now use a single streaming pass per file, enabling faster and lower-memory processing of large scan files.

Nessus STIG and CIS compliance results now generate security checklists and POA&M issues instead of vulnerabilities, with re-scans automatically closing remediated compliance findings.

CLI 6.34.65 (2026-06-17) BLOG ↗ STABLE

RegScale CLI 6.34.65 changes compliance scan default behavior and documents vulnerabilityCreation options.

└──▷ **USE IT**

Opt back in to auto-creating POAMs from failed compliance control assessments, restoring pre-6.34.65 behavior.

```yaml
complianceCreation: POAM
```

Changes compliance scan default: failed control assessments no longer auto-create issues; failed controls remain visible on assessments, with opt-in via `complianceCreation` set to `Issue` or `POAM`.

Documents `vulnerabilityCreation` configuration options for controlling how vulnerability findings are handled.

CLI 6.34.59 (2026-06-16) BLOG ↗ STABLE

RegScale CLI 6.34.59 adds eMASS API sync commands and automatic artifact upload on XML import.

New `emass_api` commands enable synchronization of POA&Ms, controls, milestones, and artifacts directly with eMASS via its API.

eMASS XML import now automatically uploads referenced artifact files to the System Security Plan when those files are present alongside the export.

Supports per-asset Issue/POAM creation, allowing findings to be tracked at the individual asset level.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro

Feedback ⓘ Claude Sonnet 5 Now Available

8 RELEASES · SEEN 2026-07-13 NOTES ↗ LATEST STABLE

Kiro is an AI-powered IDE by AWS that uses "specs" to turn requirements into structured plans, then automatically generates and implements code.

Claude Sonnet 5 is now available in Kiro IDE, CLI, and Web with 1M context and stronger agentic reasoning.



└──▷ HOW TO FIND IT

Switch to Claude Sonnet 5 in the Kiro IDE or CLI to get stronger autonomous coding and reasoning for long-running agentic tasks.

📍 Restart your IDE or CLI, then open the model selector and choose Claude Sonnet 5.

Adds Claude Sonnet 5 to the Kiro IDE, CLI, and Web model selector — Anthropic's most agentic Sonnet model with stronger reasoning, tool use, and coding performance approaching Opus 4.8 at Sonnet-class pricing.

Supports a full 1M token context window with a 1.3× credit multiplier, matching Sonnet 4.6 pricing.

Available to Kiro Pro, Pro+, Pro Max, and Power customers in AWS US-East-1 (Northern Virginia) and AWS Europe (Frankfurt) with cross-region inference support.

▾ **7 more releases in this issue** · seen 2026-07-13

Config Hot-Reload and Resource Inheritance Control NOTES ↗ STABLE

Kiro adds live hot-reload for MCP/agent configs and a toggle to block default resource inheritance in custom agents.

└──▷ **USE IT**

Prevent a custom agent from pulling in global steering files and AGENTS.md when you want a fully isolated, self-contained agent context.

```yaml
chat.disableInheritingDefaultResources: true
```

Adds hot-reload for agent and MCP configurations: saving changes to an agent config, a new agent file, or mcp.json takes effect immediately without restarting the session or losing conversation history.

Adds selective server restart on config change — only affected MCP servers restart, and the config diff is order-independent so reordering environment variables does not trigger unnecessary restarts.

New `chat.disableInheritingDefaultResources` setting lets custom agents opt out of automatically inheriting default steering, skills, and AGENTS.md.

Custom Overage Caps via Service Quotas NOTES ↗ STABLE

Enterprise admins can now set per-profile overage caps for Kiro users via AWS Service Quotas.



└──▷ **HOW TO FIND IT**

Cap per-user overage spend for a Kiro profile to prevent runaway costs across your organization.

📍 In the Kiro console, go to Settings › Overages and click the Manage link — this navigates to the AWS Service Quotas console where you can set the `Maximum allowed overage per Kiro profile` quota value.

`Expanded MCP OAuth Support` **NOTES ↗** STABLE

Kiro MCP OAuth now supports pre-registered apps with client secrets, custom redirect URIs, and DCR bypass.

└──▷ **USE IT**

Connect Kiro to a remote MCP server (like Figma) that requires a pre-registered OAuth app with a client secret and custom callback URL.

```json
{
  "mcpServers": {
    "figma": {
      "url": "https://mcp.figma.com",
      "oauth": {
        "clientId": "<your-client-id>",
        "clientSecret": "<your-client-secret>",
        "redirectUri": "http://localhost:7778/oauth/callback"
      }
    }
  }
}
```

Supports configuring a `clientSecret` in MCP OAuth config for servers that require one (e.g., Figma).

Skips Dynamic Client Registration automatically when a custom `clientId` is provided, authenticating as your pre-registered app.

Accepts full URLs with custom paths in the `redirectUri` field (e.g., `http://localhost:7778/oauth/callback`) to match exact app registration requirements.

Adds full ASCII mode coverage for terminals without Unicode support.

Delivers more accurate approval prompts for commands with combined flags.

Hooks on Agent Writes, Lazy MCP Auth, and Multi-Window Sync NOTES ↗ STABLE

Kiro adds hooks on agent file writes, lazy MCP OAuth, and multi-window settings sync.

Adds hooks that trigger on agent-driven file writes, enabling automated workflows whenever the AI agent modifies files.

Introduces lazy MCP authentication — OAuth sign-in for MCP servers is triggered automatically on demand rather than requiring manual setup upfront.

Enables multi-window settings sync so IDE preferences stay consistent across multiple open Kiro windows.

IAM Roles and Authorize Powers for Third-Party Services NOTES ↗ STABLE

Kiro adds AWS IAM role assumption for sandbox tasks and a standardized OAuth flow for third-party service Powers.

└──▷ **HOW TO FIND IT**

Grant the Kiro sandbox scoped AWS access so agent tasks and MCP servers can call AWS APIs without embedding long-lived keys.

> 📍 In the console, go to Settings › Agent › Sandbox › AWS Credentials and enter the IAM role ARN that Kiro Web should assume.

New IAM Role for Sandbox: configure an AWS IAM role that Kiro Web assumes to supply the agent, CLI tools, and MCP servers with auto-refreshing short-lived credentials scoped to each task.

New standardized authorization flow for Powers: connect third-party services (Figma, Stripe, Supabase, etc.) once via Manage Powers and the agent reuses the encrypted token across all future tasks without re-prompting.

Instant revocation for third-party service connections: disconnecting a Power immediately revokes access, with tokens encrypted at rest and never exposed to the sandbox or agent.

Session Restore, Performance Optimizations, and Custom Agent Enhancements

NOTES ↗ STABLE

Kiro 1.0.89 adds automatic session restore on launch, idle performance improvements, and prepaid credit packs in the account dashboard.

Reduces idle resource consumption for improved background performance.

Adds prepaid credit packs visibility to the account dashboard.

Stay Connected to MCP Servers NOTES ↗ STABLE

Kiro adds MCP OAuth management commands and keyboard shortcuts to control remote server authentication without restarting sessions.

└──▷ TRY IT

When a remote MCP server token has gone stale mid-session, force re-authentication without restarting.

```
$ /mcp auth
```

When an OAuth browser confirmation flow hangs and blocks your session, abort it immediately.

```
$ /mcp cancel-auth
```

When rotating credentials or revoking access to a remote MCP server, clear all stored OAuth credentials.

```
$ /mcp logout
```

New `/mcp auth` command forces re-authentication when an OAuth token has expired or become invalid on a remote MCP server.

New `/mcp cancel-auth` command aborts a stuck pending OAuth flow waiting for browser confirmation.

New `/mcp logout` command removes stored OAuth credentials for a remote MCP server entirely.

New keyboard shortcuts in the MCP panel status view: `^A` to force auth, `^X` to abort, and `^R` to remove credentials.

Updated `/usage` display reflects the new prepaid credits model.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

Hermes

Feedback ⓘ v2026.7.7 3 RELEASES · 2026-06-19 → 2026-07-08

NOTES ↗ CODE ↗ LATEST STABLE

The agent that grows with you

Hermes v2026.7.7 adds session bulk-archive, expanded prune filters, 1Password secrets, one-shot usage reports, and ~80% faster first-token latency.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.7.7 https://github.com/NousResearch/hermes-agent.git
# already have the repo? check out this version:
$ git checkout v2026.7.7
```

└──▷ TRY IT

Bulk-archive recent CI smoke-test sessions by source without deleting them, then confirm what was hidden.

```
$ hermes sessions archive --source cron --started-after 2026-07-01
```

Capture a JSON token/cost usage report after a one-shot automation run for billing or audit purposes.

```
$ hermes -z "summarize /var/log/auth.log" --usage-file usage.json
```

Adds `hermes sessions archive` subcommand to bulk soft-hide sessions without deleting them.

Expands `hermes sessions prune` with a full filter surface: time bounds (`--started-before` / `--started-after`), model, provider, user/chat origin, git branch, token count, cost, tool calls, and more — filters now match all ages when any attribute filter is supplied, dropping the implicit 90-day cutoff.

Prune dry-run preview now reports match count plus the oldest and newest matching session's age span.

Adds pluggable `SecretSource` interface with a multi-source orchestrator, enabling secrets to be resolved from multiple backends in priority order.

Adds 1Password (`op://`) as a native secret source, letting Hermes resolve credentials directly from 1Password vaults.

+2 more

▾ **2 more releases in this issue** · 2026-06-19 → 2026-07-08

v2026.7.1 NOTES ↗ CODE ↗ STABLE

Hermes v2026.7.1 adds MoA as a first-class model, `/learn`, `/journey`, `/prompt`, background subagent fan-out, Vertex AI, and scale-to-zero gateway.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2026.7.1
https://github.com/NousResearch/hermes-agent.git
# already have the repo? check out this version:
$ git checkout v2026.7.1
```

└──▷ **TRY IT**

Run your hardest prompt through a named MoA ensemble to see each frontier model's reasoning before the aggregator synthesizes a final answer.

```
$ /moa
```

Teach Hermes a repeatable workflow from a directory you just walked it through, so it's available as a skill next time without manual authoring.

```
$ /learn ./scripts/deploy-workflow
```

Inspect and prune everything Hermes knows about you — edit or delete any memory or skill from the timeline.

```
$ /journey
```

Adds Mixture-of-Agents (MoA) presets as selectable virtual models under a `moa` provider, available in every model picker across CLI, TUI, desktop, and gateway.

New `/moa` command runs a single prompt through the default MoA preset as a one-shot, then restores the previous model.

Streams each reference model's full reasoning as labelled blocks before the aggregator's live-streamed synthesis, making ensemble deliberation visible.

New `/learn <anything>` command distills a reusable skill from a directory, URL, or recent workflow and writes it to the project's CONTRIBUTING.md automatically.

New `/journey` command in CLI and TUI displays a browsable, editable timeline of all memories and skills Hermes has accumulated, with delete support.

+9 more

v2026.6.19 NOTES ↗ CODE ↗ STABLE

Hermes v0.17.0 adds iMessage/Raft/WhatsApp channels, background subagents, image editing, Automation Blueprints, and a major memory tool upgrade.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.6.19
https://github.com/NousResearch/hermes-agent.git
# already have the repo? check out this version:
$ git checkout v2026.6.19
```

└──▷ TRY IT

Connect Hermes to iMessage through Photon's managed line pool — no Mac relay needed.

```
$ hermes photon login
```

Kick off a long research or build task in the background and keep the main session free while it runs.

```python
delegate_task(background=true)
```

Opt in to LLM-powered skill consolidation while keeping routine curation token-free on all other runs.

```
$ hermes curator run --consolidate
```

Adds iMessage support via Photon Spectrum — no Mac relay or BlueBubbles bridge required; authenticate with `hermes photon login` and a device code.

Adds Raft agent network integration as a new gateway channel, letting external Raft agents wake Hermes via a wake-channel bridge with metadata-only payloads.

Adds official WhatsApp Business Cloud API adapter — Meta's first-party hosted path

with no QR-scanning bridge process.

Enables background/async subagents via delegate_task(background=true), returning a handle immediately so work continues in parallel.

Extends `image_generate` with image-to-image editing — pass a source image and prompt to edit, transform, or recompose it across all supported image providers.

+13 more

WAS THIS USEFUL? 👍 👎

OpenClaw

Feedback ⓘ v2026.6.10 2 RELEASES · 2026-06-21 → 2026-06-24

NOTES ↗ CODE ↗ LATEST STABLE

Your own personal AI assistant. Any OS.

OpenClaw v2026.6.10 adds `/fast` auto mode, expanded GLM-5.2 reasoning, new catalog models, and a session-transcript SDK API for plugins.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.6.10
https://github.com/openclaw/openclaw.git
# already have the repo? check out this version:
$ git checkout v2026.6.10
```

└──▷ TRY IT

Enable automatic fast mode so quick back-and-forth replies start immediately while heavier reasoning falls back to normal mode transparently.

```
$ /fast auto
```

Run a high-reasoning agent task with the newly catalogued Kimi K2.7 Code model via OpenCode Go from the CLI.

```
$ openclaw agent --message "Review this diff for privilege escalation
risks" --thinking high
```

Check gateway configuration health (core checks now appear before extension checks for cleaner output).

```
$ openclaw doctor
```

Adds `/fast auto` command to start short conversational calls quickly, then return longer or fallback work to normal mode while keeping effective state visible.

Shows the effective automatic fast-mode state in status (not just on/off), including through fallback transitions and connected-agent sessions.

Expands `zai/glm-5.2` thinking choices beyond binary on/off — high and max requests now map to the intended Z.AI reasoning effort.

Adds GLM-5.2 and Kimi K2.7 Code to the OpenCode Go model catalog so both can be selected from OpenClaw.

Adds a durable session-transcript SDK contract so plugins can read, append, publish, and lock the intended transcript without relying on legacy file paths as identity.

+2 more

▾ **1 more release in this issue** · 2026-06-21 → 2026-06-24

v2026.6.9 NOTES ↗ CODE ↗ STABLE

OpenClaw v2026.6.9 adds Codex Hosted Search, auto plugin approvals, OpenTelemetry log export, standalone provider npm packages, and richer Telegram delivery.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2026.6.9
https://github.com/openclaw/openclaw.git
# already have the repo? check out this version:
$ git checkout v2026.6.9
```

└──▷ **TRY IT**

Install the StepFun provider plugin directly from npm so it is auto-discovered by the Gateway at next startup.

```
$ npm install -g openclaw-provider-stepfun
openclaw gateway status
```

Adds Codex Hosted Search as a new search provider available without a separate API key.

Adds automatic plugin approvals and SecretRefs for Codex integrations.

Exposes remote-node `exec` as a dynamic tool in Codex when a connected node is available.

Adds OpenTelemetry log export for observability pipelines.

Publishes official provider plugins as standalone npm packages; Gateway auto-discovers installed channel plugins at startup.

+4 more

WAS THIS USEFUL?  

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Syslifters SysReptor

Feedback ⓘ 2026.50 NOTES ↗

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

SysReptor 2026.50 moves LLM config to the web UI, adds multi-model support, and extends auto-delete to finished projects.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2026.50
https://github.com/Syslifters/sysreptor.git
# already have the repo? check out this version:
$ git checkout 2026.50
```

└──▷ USE IT

Enforce automatic deletion of finished and archived projects after 90 days — configure once in your environment rather than per-project.

```ini
AUTOMATICALLY_DELETE_PROJECTS_AFTER=90
```

Switch which LLM model is used for a session or configure available models without restarting the service.

📍 In the console, go to Settings › LLM Models and add or select from the configured `AI_AGENT_MODELS` entries.

Moves LLM configuration (`AI_AGENT_MODELS`) from environment variables into the web UI for easier management.

Supports multiple LLM models with per-session model selection directly in the UI.

Extends auto-delete policy (`AUTOMATICALLY_DELETE_PROJECTS_AFTER`) to cover both finished and archived projects, not just archived ones.

Enables customizing the deletion date on a per-project basis.

Shows note title in the browser window title for easier tab navigation.

└──▷ BREAKING ON UPGRADE

! The `AUTOMATICALLY_DELETE_ARCHIVED_PROJECTS_AFTER` configuration key is renamed to `AUTOMATICALLY_DELETE_PROJECTS_AFTER`

! LLM configuration `AI_AGENT_MODELS` is moved from environment variables to the web UI

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 13, 2026

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

▷ Subscribe Feedback on this issue